arXiv:1401.7241v2 [stat.ME] 28 Mar 2015

Markov adaptive Pólya trees and multi-resolution adaptive shrinkage in nonparametric modeling

Li Ma

March 31, 2015

Abstract

We introduce a hierarchical nonparametric model for probability measures based on a multi-resolution transformation of probability distributions. The model allows a varying amount of shrinkage to be applied to data features of different scales and/or at different locations in the sample space, and the varying shrinkage level is locally adaptive to the empirical behavior of the data. Moreover, the model's hierarchical design—through a latent Markov tree structure—allows borrowing of information across locations and scales in setting the adaptive shrinkage level. Inference under the model proceeds efficiently using general recipes for conjugate hierarchical models. We illustrate the work of the model in density estimation and evaluate its performance through simulation under several schematic scenarios carefully designed to be representative of a variety of applications. We compare its performance to those of several state-of-the-art nonparametric models—the Pólya tree, the optional Pólya tree, and the Dirichlet process mixture of normals. In addition, we establish several important theoretical properties for the model including absolute continuity, full nonparametricity, and posterior consistency.

1 Introduction

In his seminal works that jump-started modern Bayesian nonparametric inference, Ferguson (1973, 1974) formalized the notion of a Dirichlet process (DP) and introduced a tail-free process that contains the DP as a special case. This tail-free process was later named the Pólya tree (PT) due to its relationship to the Pólya urn model (Mauldin et al., 1992), and was popularized in the 1990's by a sequence of works (Lavine, 1992; Mauldin et al., 1992; Lavine, 1994) that investigated its various theoretical properties.

The PT produces probability measures through a multi-resolution generative procedure. The sample space is recursively bisected into smaller and smaller sets, and for each set A that arises during the partition, the probability assigned to A is randomly split between its two children A_l and A_r through the drawing of a Beta random variable corresponding to the proportion of mass assigned to A_l.

There is an impressive resemblance between the inference schema of the PT and that of other multi-resolution inference methods such as wavelet denoising (Draper, 1999). In particular, inference with the PT and wavelet analysis both adopt a "divide-and-conquer" strategy: a nonparametric quantity of interest is transformed into a collection of coefficients defined on a multi-resolution tree. These coefficients characterize the shape of the nonparametric quantity at different locations and scales. More specifically, for wavelets, each coefficient—called a wavelet coefficient (WC)—specifies the local contrast of the function value, while for the PT, each coefficient—which we shall refer to (and define formally later) as the probability assignment coefficient (PAC)—characterizes how probability mass is assigned locally. Estimating the underlying nonparametric quantity then proceeds through inferring the WCs and PACs. Table 1 summarizes the analogy between multi-resolution density estimation using PT and wavelet-based function estimation.

It is well-known in wavelet analysis that a key to effective estimation is appropriate shrinkage on the WCs to differentiate "signals" from "noise" (Donoho and Johnstone, 1994).

	Wavelet denoising	PT-type density estimation
Quantity of interest	mean function	probability density
Data	random function(s) with noise	i.i.d. observations from the density
Unit of inference	wavelet coefficient	probability assignment coefficient
Unit model	Gaussian experiment	binomial experiment

Table 1: Analogy between wavelet denoising and PT-type density estimation

Numerous shrinkage methods, both frequentist and Bayesian, have been proposed for achieving data-adaptive shrinkage for wavelet denoising (Vidakovic, 1999). We note that the same is true for multi-resolution density estimation—appropriate shrinkage on the PACs is also critical for effective inference. However, existing methods such as the standard PT allows no adaptivity in shrinkage. In particular, the PT model places independent Beta priors on the PACs, which applies a prespecified, fixed amount of shrinkage determined by the prior Beta variance to the PACs. This is analogous to placing independent Gaussian priors with fixed variances on the WCs in a wavelet analysis. The following example illustrates how the lack of adaptivity in shrinkage can result in poor inference with the PT, especially when the underlying distribution contains structures of different scales.

Example 1. We simulate 750 i.i.d data from the following mixture distribution on $[0, 1]$

$$0.1\,\mathrm{U}(0,1) + 0.3\,\mathrm{U}(0.25,0.5) + 0.4\,\mathrm{Beta}_{(0.25,0.5)}(2,2) + 0.2\,\mathrm{Beta}(6000,4000)$$

where $\mathrm{Beta}_{(0.25,0.5)}(2,2)$ represents a $\mathrm{Beta}(2,2)$ translated and scaled to be supported on the interval (0.25,0.5)—that is, the distribution with density $8(4x-1)(1-2x)$ on $(0.25, 0.5)$. Figure 1 illustrates the pdf (red dashed). The "hump" on the interval $(0.25, 0.5)$ constitutes a distributional structure of a relatively large scale or low resolution, while the spike given by $\mathrm{Beta}(6000, 4000)$ constitutes a small-scale or high-resolution feature.

Let us place a PT prior on the underlying distribution corresponding to a $\mathrm{Beta}(k^2, k^2)$ prior on for PACs at level k, which is the most common specification in applications of the PT. (See Lavine (1992); Walker et al. (1999); Hanson and Johnson (2002); Hanson (2006);

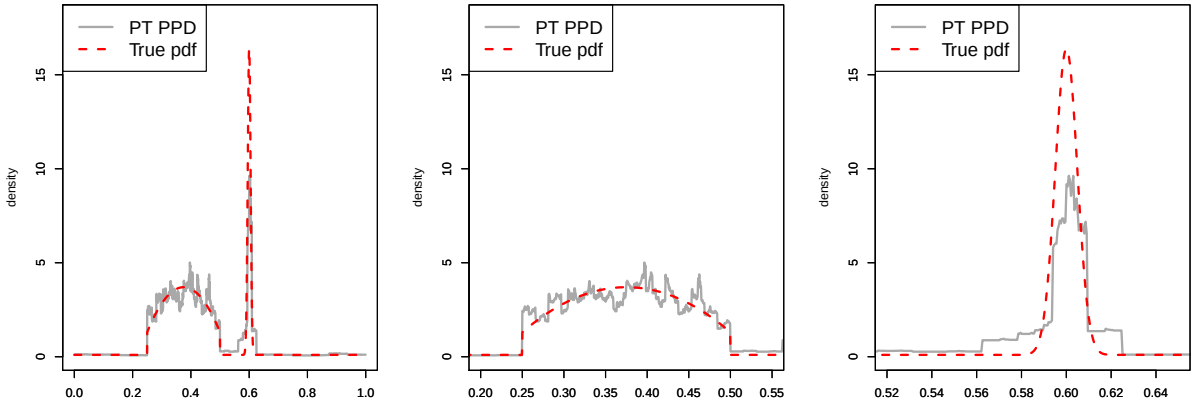

Figure 1: The true density (red dashed) and the PPD of the PT (gray solid) for Example 1. The middle and right plots give the zoom-in views of the low- and high-resolution features.

Holmes et al. (2009) for example.) The gray solid curves in Figure 1 shows the posterior predictive density (PPD) of the PT. The middle and right plots give the zoom-in views of the large-scale and small-scale features respectively. We see that overall the PT induces a decent amount of shrinkage for capturing the low-resolution feature (middle), as the general shape of that feature is adequately recovered in the PPD. For the high-resolution feature, however, it results in too much shrinkage and thus under-estimation of the mode (right). Interestingly, if we zoom into higher resolutions within the low-resolution feature (middle), the PPD shows jumpy patterns of overfitting, indicating that more shrinkage is needed there at higher resolutions to ensure proper smoothness.

The above example represents a typical situation—the appropriate amount of shrinkage varies across locations and scales. Only an oracle can *a priori* choose the prior variance of the PACs that give the optimal amount of shrinkage at all resolutions and locations. Popular convenient choices such as k^2 for the Beta parameters—or any other fixed function of k for that matter—cannot provide the right amount of shrinkage at all locations and scales, just as independent Gaussian priors with pre-specified variances for the WCs are inadequate in wavelet denoising (Chipman et al., 1997; Vidakovic, 1998; Clyde and George, 2000a,b).

Our main goal is to incorporate adaptive shrinkage into the multi-resolution nonparametric modeling framework adopted in the PT and related models, which we shall refer to generally as "PT-type" models. Both classical thresholding and Bayesian shrinkage methods can be adopted, and we shall take a hierarchical Bayesian approach focusing on constructing hyperpriors in the form of generative models on the shrinkage parameters, i.e. the prior variance of the PACs. Our modeling approach provides a natural way to introducing dependency into the shrinkage levels for different PACs, thereby achieving borrowing of information across locations and scales in determining the level of shrinkage for each PAC. In particular, we construct a joint generative model using a latent variable representation that takes the form of a Markov tree (Crouse et al., 1998) to achieve *stochastically* increasing shrinkage (defined later), which is critical for achieving effective locally adaptive smoothing in density estimation. This is in contrast to the *deterministically* increasing shrinkage imposed by the PT with prespecified increasing Beta parameters.

The rest of the work is organized as follows. In Section 2, we present the core of our methodology. We begin by viewing the PT from a hierarchical modeling viewpoint, and interpret inference under the PT from a shrinkage perspective. Then we construct a simple hyperprior on the variance of PACs that treats the PACs independently. Finally we extend the hyperprior through a latent variable representation to a Markov tree hyperprior that incorporates dependency in the shrinkage levels. We provide guidelines for prior specification and establish both methodological and theoretical properties of the model. Specifically, we show that the model is fully nonparametric (i.e. with full prior support) and enjoys posterior consistency. Moreover, posterior inference under this model can be carried out conveniently using general recipes for conjugate hierarchical models. In particular, the full posterior can be analytically derived and sampled from exactly without resorting to MCMC. In Section 3 we illustrate how our method works in density estimation and evaluate its performance under various scenarios where the underlying density possesses a variety of features. We also

compare our method to two state-of-the-art PT-type multi-resolution models—the PT and the optional Pólya tree (OPT) (Wong and Ma, 2010)—as well as to the very popular Dirichlet process mixture (DPM) of normals (Escobar and West, 1995). We conclude in Section 4 with brief remarks on applications to hypothesis testing and computational efficiency.

We close this introduction by noting that adaptive shrinkage has been extensively studied in the context of wavelet denoising from both frequentist and Bayesian perspectives. The literature is too enormous to be enumerated. A far-from-exhaustive list of notable examples from the frequentist perspective include Donoho and Johnstone (1994); Donoho (1995); Donoho and Johnstone (1995); Donoho et al. (1995); Nason and Silverman (1995); Abramovich and Benjamini (1995); Nason (1996); Johnstone and Silverman (1997); Donoho and Johnstone (1998); Cai and Brown (1998); Hall et al. (1998); Kolaczyk (1999); Johnstone (1999); Cai and Brown (1999); Antoniadis and Sapatinas (2001); Johnstone and Silverman (2005), and from the Bayesian perspective include Chipman et al. (1997); Abramovich et al. (1998); Clyde et al. (1998); Crouse et al. (1998); Vidakovic (1998); Vannucci and Corradi (1999); Moulin and Liu (1999); Chang et al. (2000); Clyde and George (2000b); Romberg et al. (2001); Brown et al. (2001); Clyde and George (2000a); Portilla et al. (2003); Morris and Carroll (2006). For the particular application of density estimation, there is also a body of literature on wavelet-based methods. See for example Vannucci (1995); Donoho et al. (1996); Koo and Kooperberg (2000); Herrick et al. (2001).

2 Method

2.1 Multi-resolution representation of probability distributions

We start by introducing some basic concepts and definitions that form the building blocks for multi-resolution modeling of probability distributions. Throughout this work, we let Ω denote the sample space, which can be finite or a (possibly unbounded) Euclidean rectangle such as an interval in $\mathbb{R}$ or a rectangle in $\mathbb{R}^p$. Let μ be the natural measure associated with

Ω, which is the counting measure if Ω is finite and the Lebesgue measure if Ω is Euclidean.

Let $\mathcal{A}^1, \mathcal{A}^2, \ldots, \mathcal{A}^k, \ldots$ be a *sequence of nested dyadic partitions* of Ω. That is, each $\mathcal{A}^k = \{A_{k,1}, A_{k,2}, \ldots, A_{k,2^k}\}$ and it satisfies (i) $\Omega = \cup_{m=1}^{2^k} A_{k,m}$, (ii) $A_{k,m_1} \cap A_{k,m_2} = \emptyset$ for all $m_1 \neq m_2$, and (iii) $A_{k,m} = A_{k+1,2m-1} \cup A_{k+1,2m}$ for all $k = 1, 2, \ldots$ and $m = 1, 2, \ldots, 2^k$. In other words, the partition $\mathcal{A}_{k+1}$ is obtained by dividing each set in $\mathcal{A}^k$ into two children, the *left child* $A_{k+1,2m-1}$ and the *right child* $A_{k+1,2m}$. We shall call $\mathcal{A}^k$ the partition at resolution (or scale) k. Also, we let $\mathcal{A}^{(\infty)} = \cup_{k=1}^{\infty} \mathcal{A}^k$, the totality of all partition sets that arise in all resolution levels. The partition sets form a bifurcating tree, so from now on we shall refer to $\mathcal{A}^{(\infty)}$ as the *partition tree*, and each A in $\mathcal{A}^{(\infty)}$ as a *node*. Because each A corresponds to a particular location and scale in the tree, we also interchangeably refer to a node as a *location-scale combination*.

Given a partition tree $\mathcal{A}^{(\infty)}$ that generates the Borel σ-algebra, one can describe a probability distribution G by specifying how probability mass is split between the left and right children on each node A. Let A_l and A_r be the left and right children of a node A. We define the *probability assignment coefficient* (PAC) for A to be the proportion of probability mass assigned to A_l, and denote it as $\theta(A)$. So if the total probability mass on A is $G(A)$ then those assigned to the children are $G(A_l) = G(A)\theta(A)$ and $G(A_r) = G(A)(1 - \theta(A))$.

Lemma 1. *Given a partition tree $\mathcal{A}^{(\infty)}$ that generates the Borel σ-algebra, every probability distribution G can be mapped to a collection of PACs $\{\theta(A) : A \in \mathcal{A}^{(\infty)}\}$, and the mapping is unique on all As such that $G(A) > 0$.*

Remark: A collection of PACs corresponding to G is given by $\theta(A) = G(A_l)/G(A)$ for all A with $G(A) > 0$ and $\theta(A) = 0$ otherwise.

Figure 2(a) illustrates the transformation of a distribution into PACs. Each PAC specifies the structure of the distribution at a given scale and location.

The lemma implies that inference on a distribution can be achieved by inferring the PACs, which motivates a "divide-and-conquer" strategy for nonparametric inference. More

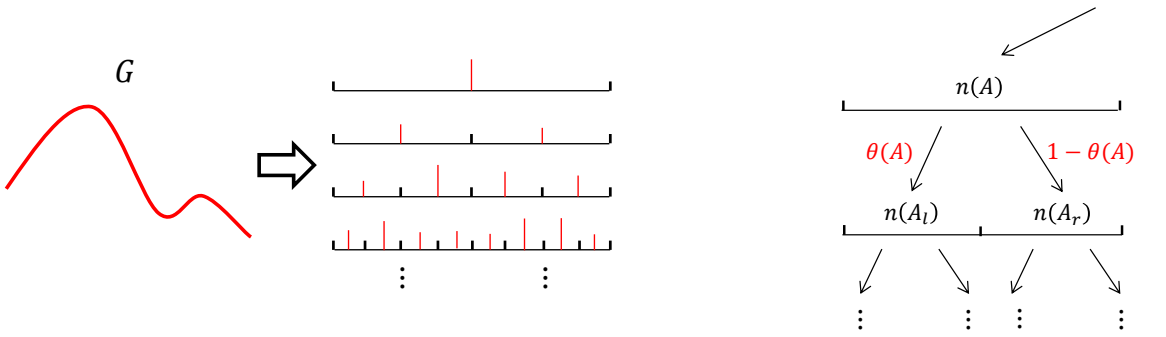

(a) Transforming a distribution into PACs (red ticks)

(b) Local binomial experiment

Figure 2: The divide-and-conquer inference schema.

specifically, the multi-resolution transformation of G into PACs induces a corresponding decomposition of the statistical experiment that generates an i.i.d. sample of size n from G. In particular, the experiment is divided into a collection of local binomial experiments carried out sequentially in the order of the resolution: $\{n(A_l) \sim \text{Binomial}\,(n(A), \theta(A)) : A \in \mathcal{A}^k\}$ for $k = 1, 2, \ldots$ where $n(A)$ is the number of data points in A arising from the binomial experiment on A's parent in the previous resolution $k-1$, except that for $k = 1$, $n(A) = n$ by design. See Figure 2(b) for an illustration of the local binomial experiment.

Accordingly, inferring a distribution through the PACs is divided into inference on the success probabilities of a collection of sequential binomial experiments. Viewed this way, the PT model provides a simple solution to this problem—it places independent conjugate Beta priors on the successes probabilities. The posterior conjugacy of the PT follows immediately from the Beta-binomial conjugacy. Thus inference under PT model is analytically tractable and computationally efficient.

2.2 Multi-resolution shrinkage and the adaptive Pólya tree

One can now understand the shrinkage property of the PT by viewing each of the binomial experiment from a shrinkage perspective. Under the PT model, $\theta(A) \sim \text{Beta}(\alpha_l(A), \alpha_r(A))$ for all $A \in \mathcal{A}^{(\infty)}$. (A popular specification has $\alpha_l(A) = \alpha_r(A) = k^2$ for $A \in \mathcal{A}^k$.) We shall

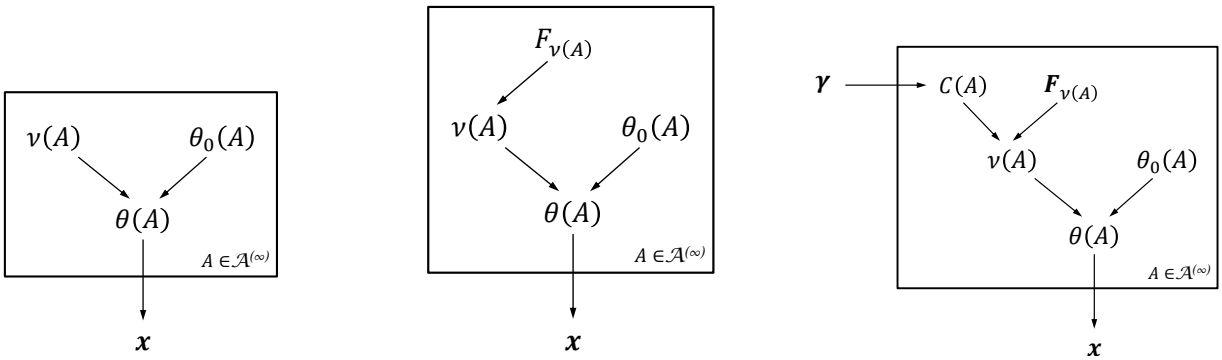

(a) Pólya tree (b) Adaptive Pólya tree (c) Markov adaptive Pólya tree

Figure 3: Graphical representation of PT (left), APT (middle), and Markov-APT (right).

prefer an alternative parametrization of Beta distributions in terms of a mean parameter $\theta_0(A) = \alpha_l(A)/(\alpha_l(A) + \alpha_r(A))$ and a precision parameter $\nu(A) = \alpha_l(A) + \alpha_r(A)$. The posterior distribution of $\theta(A)$ is still Beta with mean parameter

$$\tilde{\theta}_0(A) = E(\theta(A)|\boldsymbol{x}) = \theta_0(A) \cdot \frac{\nu(A)}{\nu(A) + n(A)} + \frac{n(A_l)}{n(A)} \cdot \frac{n(A)}{\nu(A) + n(A)}$$

and precision parameter $\tilde{\nu}(A) = \nu(A) + n(A)$. The posterior mean is a weighted average between $\theta_0(A)$, or the prior mean, and $n(A_l)/n(A)$, or the PAC on A of the empirical distribution. The level of shrinkage for $\theta(A)$ is controlled by the precision parameter $\nu(A)$, and thus we shall refer to $\nu(A)$ also as the *shrinkage parameter*.

The prior mean of the PT is the probability distribution corresponding to the collection of PACs $\{\theta_0(A) : A \in \mathcal{A}^{(\infty)}\}$, which we call Q_0, while its posterior mean is the distribution corresponding to $\{\tilde{\theta}_0(A) : A \in \mathcal{A}^{(\infty)}\}$ as the PACs. Intuitively, the posterior mean of a PT is a weighted average between Q_0 and the empirical distribution, but the weighting is scale and location dependent. Figure 3(a) provides a graphical model representation of the PT.

From a hierarchical Bayesian perspective, we can incorporate adaptivity into the shrinkage by placing a hyperprior $F_{\nu(A)}$ on $\nu(A)$ thereby allowing the appropriate level of shrinkage for $\theta(A)$ to be inferred from the data. This leads to a generative hierarchical model specified by $F_{\boldsymbol{\nu}} = \{F_{\nu(A)} : A \in \mathcal{A}^{(\infty)}\}$ and Q_0. Using $\boldsymbol{\phi} = (F_{\boldsymbol{\nu}}, Q_0)$ to represent the totality of all

hyperparameters, we can write the model as

$$\nu(A) \,|\, \boldsymbol{\phi} \sim F_{\nu(A)}$$

$$\theta(A) \,|\, \boldsymbol{\phi}, \boldsymbol{\nu} \sim \text{Beta}(\theta_0(A)\nu(A), (1-\theta_0(A))\nu(A))$$

for all $A \in \mathcal{A}^{(\infty)}$, where $\boldsymbol{\nu} = \{\nu(A) : A \in \mathcal{A}^{(\infty)}\}$, the collection of the Beta precision parameters. Figure 3(b) provides a graphical representation of this model.

We consider priors $F_{\nu(A)}$ supported on $(0, \infty]$. Note that we allow $\nu(A) = \infty$, in which case $\text{Beta}(\theta_0(A)\nu(A), (1-\theta_0(A))\nu(A))$ is a point mass at $\theta_0(A)$, corresponding to complete shrinkage of $\theta(A)$ to the prior mean $\theta_0(A)$. The first natural question is whether this model always generates well-defined probability measures. In other words, given a set of PACs $\{\theta(A) : A \in \mathcal{A}^{(\infty)}\}$ arising from the above model, does there (almost surely) exist a distribution G such that $G(A_l|A) = \theta(A)$ for all $A \in \mathcal{A}^{(\infty)}$? The answer is positive by Theorem 3.3.2 in Ghosh and Ramamoorthi (2003). Hence we can define this model as a distribution on probability measures.

Definition 1 (Adaptive Pólya tree)**.** A probability measure Q is said to have an *adaptive Pólya tree* (APT) distribution with parameters $\boldsymbol{\phi} = (\boldsymbol{F_\nu}, Q_0)$ if the corresponding PACs of Q, $\{\theta(A) : A \in \mathcal{A}^{(\infty)}\}$, are generated from the above hierarchical model. We write $Q \sim \text{APT}(\boldsymbol{F_\nu}, Q_0)$, or equivalently $\text{APT}(\boldsymbol{F_\nu}, \boldsymbol{\theta}_0)$.

Remark: A hidden hyperparameter is the partition tree $\mathcal{A}^{(\infty)}$, which for simplicity we treat as given. One can also treat it as a parameter and even place a further layer of hyperprior on $\mathcal{A}^{(\infty)}$ as in Hanson (2006), resulting in a class of mixture of APTs.

The meaning of Q_0 stays the same as for the PT—it is still the mean of the new model.

Theorem 1 (Mean)**.** *The mean of an APT distribution is Q_0. That is, for any Borel set $B \subset \Omega$, a random measure Q that has the APT distribution satisfies $EQ(B) = Q_0(B)$.*

Under very general prior specifications, the APT model has full L_1 support and enjoys

posterior consistency. We defer the study of these properties to Section 2.3 where we derive these results for a more general model that contains the APT as a special case. The next theorem provides the Bayesian inference recipe for the APT.

Theorem 2 (Posterior)**.** *Suppose* $Q \sim \mathrm{APT}(\boldsymbol{F_\nu}, Q_0)$*. Given* n *i.i.d. observations* $\boldsymbol{x} = (x_1, x_2, \ldots, x_n)$ *from* Q*, the joint posterior of* $(\boldsymbol{\nu}, \boldsymbol{\theta})$ *is given by*

$$\begin{aligned} \nu(A) \,|\, \boldsymbol{\phi}, \boldsymbol{x} &\sim \tilde{F}_{\nu(A)} \\ \theta(A) \,|\, \boldsymbol{\phi}, \boldsymbol{\nu}, \boldsymbol{x} &\sim \mathrm{Beta}(\tilde{\theta}_0(A)\tilde{\nu}(A), (1-\tilde{\theta}_0(A))\tilde{\nu}(A)) \end{aligned}$$

where $\tilde{\theta}_0(A)$ *and* $\tilde{\nu}(A)$ *are as defined before, while* $\tilde{F}_{\nu(A)}$ *is given by*

$$d\tilde{F}_{\nu(A)}(\nu) = dF_{\nu(A)}(\nu) \cdot M_A(\boldsymbol{\theta}_0, \nu)/M_A(\boldsymbol{\theta}_0)$$

with

$$M_A(\boldsymbol{\theta}_0, \nu) = \frac{\Gamma(\theta_0(A)\nu + n(A_l))\Gamma((1-\theta_0(A))\nu + n(A_r))\Gamma(\nu)}{\Gamma(\nu + n(A))\Gamma(\theta_0(A)\nu)\Gamma((1-\theta_0(A))\nu)}$$

and

$$M_A(\boldsymbol{\theta}_0) = \int M_A(\boldsymbol{\theta}_0, \nu) dF_{\nu(A)}(\nu).$$

Remark: $M_A(\boldsymbol{\theta}_0, \nu)$ is the marginal likelihood of the local binomial experiment on A given $\boldsymbol{\theta}_0$—in particular $\theta_0(A)$—and $\nu(A) = \nu$, while $M_A(\boldsymbol{\theta}_0)$ is the marginal likelihood of the local binomial experiment given $\theta_0(A)$. Note that when $\nu = \infty$, $M_A(\boldsymbol{\theta}_0, \infty) := \lim_{\nu \uparrow \infty} M_A(\boldsymbol{\theta}_0, \nu) = \theta_0(A)^{n(A_l)}(1-\theta_0(A))^{n(A_r)}$, which is still the likelihood of the local binomial experiment.

Because ν is one-dimensional, it is easy to evaluate $M_A(\boldsymbol{\theta}_0)$ numerically. Specifically, we evaluate $f_{\nu(A)} = dF_{\nu(A)}/d\mu$ on a grid of different ν values $\nu_{(1)}, \nu_{(2)}, \ldots, \nu_{(H)}$ covering the support of $F_{\nu(A)}$, and approximate $M_A(\boldsymbol{\theta}_0)$ using a finite Riemann integral. This approximation becomes particularly straightforward when $F_{\nu(A)}$ is uniform on an interval $[a, b]$ under some transformed scale of ν, such as $\log_{10}(\nu)$. In this case one can choose the grid points to be of equal distance on the transformed scale and $M_A(\boldsymbol{\theta}_0) \approx \frac{1}{H}\sum_{h=1}^{H} M_A(\boldsymbol{\theta}_0, \nu_{(h)})$.

2.3 Stochastically increasing shrinkage and the Markov-APT

Under the APT, the $\nu(A)$'s are mutually independent both *a priori* and *a posteriori*. The shrinkage level for each PAC is inferred based on just the empirical evidence from the corresponding local binomial experiment. Our next objective is to allow borrowing of information across the binomial experiments in determining the shrinkage level for each.

What information can and should be borrowed across the experiments? The answer very much depends on the inference problem at hand. In density estimation, a reasonable assumption adopted (explicitly or implicitly) in all statistical methods is the smoothness of the underlying density. Indeed, even "jumpy" densities—those with sharp changes—must be assumed to eventually smooth out at high enough resolution (or at a small enough "bandwidth") as opposed to infinitely oscillating in arbitrarily small regions, because otherwise reliable estimation is infeasible. In the multi-resolution framework, smoothness translates into an increase in shrinkage for higher resolutions. This is exactly the motivation for the increasing Beta parameters in the PT (Lavine, 1992).

One approach to incorporating increasing shrinkage with resolution in the APT is to choose the prior $F_{\nu(A)}$ such that it is supported on larger values for nodes in deeper levels of the partition tree. For example, one may let $F_{\nu(A)}$ be supported on $[l(k), u(k))]$ for $A \in \mathcal{A}^k$ while $l(k)$ and $u(k)$ are fixed increasing functions in k. However, this strategy will inherit the very same issues of the standard PT as illustrated in the Example 1. First, the proper rate at which shrinkage should increase—in terms of the choice of $l(k)$ and $u(k)$—is typically unknown *a priori*. Second, in many applications, the smoothness of the underlying density is *heterogeneous* across the sample space—sharp boundaries or sudden jumps may lie in an otherwise smooth density. Such features are particularly common in applications involving anomaly or change-point detection, in which the jumps and boundaries are the main focus of inference. To capture such structures, one need the rate at which shrinkage increases with the resolution to vary across the sample space, which is again impossible to specify *a priori*.

Fortunately, one can infer the proper (varying) rate of increasing shrinkage from the data What we need is a *stochastic* model for the increasing rate of shrinkage. To this end, we first introduce a latent mixture representation for $F_{\nu(A)}$. Specifically, we specify $F_{\nu(A)}$ using a mixture of I component distributions in a monotone increasing stochastic order

$$F^1_{\nu(A)} \prec F^2_{\nu(A)} \prec \ldots \prec F^I_{\nu(A)}.$$

In particular, we choose these components to have non-overlapping supports. For example $F^i_{\nu(A)}$ may be supported on an interval $(a(i), b(i)]$ where $0 \leq a(1) < b(1) \leq a(2) < b(2) \ldots \leq a(I) < b(I) \leq \infty$ is a sequence of increasing support boundary points. For now we shall treat the number of components I and each $F^i_{\nu(A)}$ as given, but will provide guidelines on choosing them in Section 2.5. We let $\boldsymbol{F}_{\nu(A)} = (F^1_{\nu(A)}, F^2_{\nu(A)}, \ldots, F^I_{\nu(A)})$, and use $\boldsymbol{F}_{\boldsymbol{\nu}} = \{\boldsymbol{F}_{\nu(A)} : A \in \mathcal{A}^{(\infty)}\}$ to denote the totality of all component distributions.

Moreover, we introduce a latent state variable $C(A)$ for each $A \in \mathcal{A}^{(\infty)}$ that indicates the mixture component $\nu(A)$ comes from:

$$\nu(A) \,|\, C(A) = i \;\sim\; F^i_{\nu(A)}.$$

for $i = 1, 2, \ldots, I$. We refer to $C(A)$ as the *shrinkage state* on A, and let $\mathcal{C} = \{C(A) : A \in \mathcal{A}^{(\infty)}\}$ be the collection of all shrinkage states.

Now we can enforce stochastically non-decreasing shrinkage along each branch of $\mathcal{A}^{(\infty)}$ by specifying a joint prior on $\mathcal{C}$ that prevents the shrinkage state from moving lower in any branch. That is, if A_p is A's parent in $\mathcal{A}^{(\infty)}$, then we require $C(A) \geq C(A_p)$. A simple stochastic model for $\mathcal{C}$ that can help us impose such a constraint is the Markov tree (MT) (Crouse et al., 1998), which links the $C(A)$'s using a Markov process such that the shrinkage state $C(A)$ depends on that of A_p through Markov transition. The Markov process is initiated on the root, Ω, as follows

$$P(C(\Omega) = i) = \gamma_i(\Omega) \quad \text{for } i \in \{1, 2, \ldots, I\},$$

where the $\gamma_i(\Omega)$'s are called the *initial state probabilities*, and can be put into a vector

$$\boldsymbol{\gamma}(\Omega) = (\gamma_1(\Omega), \gamma_2(\Omega), \ldots, \gamma_I(\Omega))$$

Then for each $A \neq \Omega$, $C(A)$ is determined sequentially based on its parent according to

$$P(C(A) = i' \,|\, C(A_p) = i) = \gamma_{i,i'}(A) \quad \text{for } i, i' \in \{1, 2, \ldots, I\}$$

where $\gamma_{i,i'}(A)$ is called the *state transition probability*, which can be organized into a *transition probability matrix*

$$\boldsymbol{\gamma}(A) = \begin{pmatrix} \gamma_{1,1}(A) & \gamma_{1,2}(A) & \cdots & \gamma_{1,I}(A) \\ \gamma_{2,1}(A) & \gamma_{2,2}(A) & \cdots & \gamma_{2,I}(A) \\ \vdots & \vdots & \vdots & \vdots \\ \gamma_{I,1}(A) & \gamma_{I,2}(A) & \cdots & \gamma_{I,I}(A) \end{pmatrix}.$$

The desired stochastically increasing shrinkage is achieved when the transition matrices are all upper-triangular. That is,

$$\gamma_{i,i'}(A) = 0 \quad \text{if } i > i' \text{ for all } A \in \mathcal{A}^{(\infty)}.$$

From now on, we shall use $\boldsymbol{\gamma} = \{\boldsymbol{\gamma}(A) : A \in \mathcal{A}^{(\infty)}\}$ to denote the collection of all initial state probabilities and transition probability matrices needed for specifying the MT.

Putting the pieces together, now we have the following hierarchical model for a probability distribution with hyperparameters $\boldsymbol{\phi} = (\boldsymbol{\gamma}, \boldsymbol{F_\nu}, Q_0)$

$$\begin{aligned} \mathcal{C} \,|\, \boldsymbol{\phi} &\sim \mathrm{MT}(\boldsymbol{\gamma}) \\ \nu(A) \,|\, \boldsymbol{\phi}, \mathcal{C} &\sim \sum_{c=1}^{I} F^i_{\nu(A)} \cdot \mathbf{1}_{C(A)=i} \\ \theta(A) \,|\, \boldsymbol{\phi}, \boldsymbol{\nu}, \mathcal{C} &\sim \mathrm{Beta}(\theta_0(A)\nu(A), (1-\theta_0(A))\nu(A)) \end{aligned}$$

for all $A \in \mathcal{A}^{(\infty)}$. A graphical representation of this model is given in Figure 3(c). Because this hierarchical model also generates a probability measure with probability 1, one can again define it formally as a distribution on probability measures.

Definition 2 (Markov adaptive Pólya tree)**.** A probability measure Q is said to have a *Markov adaptive Pólya tree* (Markov-APT) distribution with parameters $\boldsymbol{\phi} = (\boldsymbol{\gamma}, \boldsymbol{F_\nu}, Q_0)$ if Q corresponds to the collection of PACs $\{\theta(A) : A \in \mathcal{A}^{(\infty)}\}$ generated from the above hierarchical model. We write $Q \sim \text{Markov-APT}(\boldsymbol{\gamma}, \boldsymbol{F_\nu}, Q_0)$. In particular, when $\boldsymbol{\gamma}(A)$ is upper-triangular for all $A \in \mathcal{A}^{(\infty)}$, we say that the Markov-APT is *stochastically increasing*.

The meaning of Q_0 is still the same as before.

Theorem 3 (Mean)**.** *The mean of an Markov-APT is Q_0. That is, for any Borel set $B \subset \Omega$, a random measure Q that has a Markov-APT distribution satisfies $EQ(B) = Q_0(B)$.*

In density estimation, we shall focus on stochastically increasing Markov-APTs. Moreover, in order to model a density, we must ensure that a random measure Q generated from an Markov-APT has a density. Earlier works in the literature have established two general approaches for achieving absolute continuity for PT-type priors, which could both be adopted for the Markov-APT. The first is to force $\nu(A)$ to increase with the level of A at a sufficiently fast rate (Lavine, 1992). For the Markov-APT, this can be achieved by choosing $\boldsymbol{F}_{\nu(A)}$ such that for any $A \in \mathcal{A}^k$, $\nu(A) > l(k)$ with probability 1 where $l(k)$ is a positive function in k that satisfies $\sum_{k=1}^{\infty} 1/l(k) < \infty$ (Kraft, 1964). But this would impose a minimum amount of prespecified shrinkage homogeneously across the sample space, which is exactly the undesirable feature of the PT that we wish to avoid through adaptive shrinkage. For this reason we prefer an alternative strategy for ensuring absolute continuity (Wong and Ma, 2010), which is to include a separate "complete shrinkage" state as in the following theorem.

Theorem 4 (Absolute continuity)**.** *Suppose Q has a stochastically increasing Markov-APT distribution for which $F^I_{\nu(A)} = \mathbf{1}_\infty$, a point mass at ∞, and there exists $\delta > 0$ such that for all large enough k, the state transition probabilities for any $A \in \mathcal{A}^k$ satisfies*

$$\gamma_{i,I}(A) > \delta \quad \text{for } i = 1, 2, \ldots, I-1,$$

then with probability 1, $Q \ll Q_0$. In particular, if $Q_0 \ll \mu$, then $Q \ll \mu$.

Remark: The complete shrinkage state eliminates the need for increasing lower bound on the support of $\nu(A)$ to ensure absolute continuity.

Next we establish two theoretical guarantees for inference using the Markov-APT model. The first property regards the flexibility of the model—it shows that Markov-APT enjoys full prior support. Thus inference with the Markov-APT is fully nonparametric.

Theorem 5 (Large prior support)**.** *Suppose $Q \sim$ Markov-APT$(\boldsymbol{\gamma}, \boldsymbol{F_\nu}, Q_0)$ that satisfies the conditions of Theorem 4. In addition, suppose (i) $Q_0 \ll \mu$, (ii) $I \geq 2$, (iii) $\exists \delta' > 0$ such that for all large enough k, $\gamma_{i,I}(A) < 1 - \delta'$ for all $A \in \mathcal{A}^k$ and $i = 1, 2, \ldots, I-1$, and (iv) $\exists \epsilon > 0$ and $N > 0$ such that $F^i_{\nu(A)}\big((0, N]\big) > \epsilon$ for all $i = 1, 2, \ldots, I-1$ and all $A \in \mathcal{A}^{(\infty)}$. Then for any distribution $G \ll Q_0$ and any $\tau > 0$, we have*

$$\pi\left(Q : \int |q - g| d\,\mu < \tau\right) > 0$$

where $q = dQ/d\mu$ and $g = dG/d\mu$ are the corresponding densities with respect to μ.

The next result regards the asymptotic consistency of inference using the Markov-APT and it guarantees that as we get more and more data, the posterior will eventually concentrate into any weak neighborhood of the true density. For any probability measures P_0 on Ω, a *weak neighborhood* U of P_0 is a set of probability measures on Ω of the form

$$U = \left\{ Q : \Big| \int f_i(\cdot) dQ - \int f_i(\cdot) dP_0 \Big| < \epsilon_i, \text{ for } i = 1, 2, \ldots, K \right\}$$

for any bounded continuous functions f_i's and non-negative constants ϵ_i's.

Theorem 6 (Posterior consistency under weak topology)**.** *Suppose $X_1, X_2, \ldots, X_n, \ldots$ are i.i.d. data from Q, and let $\pi(\cdot)$ be a stochastically increasing Markov-APT prior on Q that satisfies the conditions in Theorem 5 and let $\pi(\cdot | X_1, X_2, \ldots, X_n)$ be the corresponding posterior. Then for any $P_0 \ll Q_0$ with bounded density dP_0/dQ_0, we have*

$$\pi(U \,|\, X_1, X_2, \ldots, X_n) \longrightarrow 1 \quad \text{as } n \to \infty$$

with $P_0^{(\infty)}$ probability 1 for any weak neighborhood U of P_0.

2.4 Bayesian inference with the Markov-APT

Next we address how to carry out posterior inference for the Markov-APT. We show that the full posterior can be derived following a general recipe for hierarchical models (Gelman et al., 2013, Sec. 5.3) and inference can proceed in a usual manner through drawing a sample from the posterior and/or computing some summary statistic such as the posterior mean. In particular, the posterior is available analytically and so no Markov Chain Monte Carlo (MCMC) is needed. In particular, the full posterior $\pi(\mathcal{C}, \boldsymbol{\nu}, \boldsymbol{\theta} \,|\, \boldsymbol{\phi}, \boldsymbol{x})$ can be described in three pieces: (i) $\pi(\boldsymbol{\theta} \,|\, \boldsymbol{\phi}, \boldsymbol{\nu}, \mathcal{C}, \boldsymbol{x})$, (ii) $\pi(\boldsymbol{\nu} \,|\, \boldsymbol{\phi}, \mathcal{C}, \boldsymbol{x})$, and (iii) $\pi(\mathcal{C} \,|\, \boldsymbol{\phi}, \boldsymbol{x})$ as follows.

(i) $\pi(\boldsymbol{\theta} \,|\, \boldsymbol{\phi}, \boldsymbol{\nu}, \mathcal{C}, \boldsymbol{x})$. This conditional posterior follows directly from the Beta-binomial conjugacy. Specifically, we have

$$\theta(A) \,|\, \boldsymbol{\phi}, \boldsymbol{\nu}, \boldsymbol{x} \sim \text{Beta}(\tilde{\theta}_0(A)\tilde{\nu}(A), (1-\tilde{\theta}_0(A))\tilde{\nu}(A)) \quad \text{for all } A \in \mathcal{A}^{(\infty)},$$

where as before $\tilde{\nu}(A) = \nu(A) + n(A)$ and $\tilde{\theta}(A) = \left(\theta_0(A)\nu(A) + n(A_l)\right)/\tilde{\nu}(A)$.

(ii) $\pi(\boldsymbol{\nu} \,|\, \boldsymbol{\phi}, \mathcal{C}, \boldsymbol{x})$. Due to the conjugacy of finite mixture models, the conditional posterior for $\nu(A)$ is still an I-component mixture

$$\nu(A) \,|\, \boldsymbol{\phi}, \mathcal{C}, \boldsymbol{x} \sim \sum_{c=1}^{I} \tilde{F}^i_{\nu(A)} \cdot \mathbf{1}_{C(A)=i}$$

where each new mixture component distribution $\tilde{F}^i_{\nu(A)}$ has density

$$d\tilde{F}^i_{\nu(A)}(\nu) = dF^i_{\nu(A)}(\nu) \cdot M_A(\boldsymbol{\theta}_0, \nu)/M^i_A(\boldsymbol{\theta}_0)$$

for $i = 1, 2, \ldots, I$. The normalizing constant (or partition function) in the above density

$$M^i_A(\boldsymbol{\theta}_0) = \int M_A(\boldsymbol{\theta}_0, \nu) dF^i_{\nu(A)}(\nu)$$

is the marginal likelihood of the local binomial experiment on A given $\theta_0(A)$ and $C(A) = i$, which can be numerically evaluated as described following Theorem 2 for $M_A(\boldsymbol{\theta}_0)$.

(iii) $\pi(\mathcal{C} \mid \boldsymbol{\phi}, \boldsymbol{x})$. The last piece is the marginal posterior on $\mathcal{C}$, which follows again from the conjugacy of finite mixtures (of which the MT is a special case) to be an MT with initial and transition probabilities computable analytically through a forward-backward algorithm (Liu, 2001). The forward step, or the summation step, is a bottom-up (leaf-to-root) recursion on the partition tree and the backward step, or the sampling step, a top-down recursion.

To describe the algorithm, we first define a mapping ξ_A for each $A \in \mathcal{A}^{(\infty)}$ as follows

$$\xi_A(i, \boldsymbol{\phi}) := \begin{cases} \int q(\boldsymbol{x}|A)\pi(dq \mid \boldsymbol{\phi}, C(A_p) = i) & \text{if } A \in \mathcal{A}^{(\infty)} \backslash \{\Omega\} \\ \int q(\boldsymbol{x}|A)\pi(dq \mid \boldsymbol{\phi}) & \text{if } A = \Omega \end{cases}$$

for $i = 1, 2, \ldots, I$, where $q(\boldsymbol{x}|A) := \prod_{x \in A} \frac{q(x)}{Q(A)}$ with $q = dQ/d\mu$, and A_p is the parent of A in $\mathcal{A}^{(\infty)}$. This mapping gives the marginal likelihood of the "submodel" on A—the Markov-APT with A being the sample space—given that the shrinkage state on A_p is i. Note that when $A = \Omega$, it does not have a parent, and $\xi_\Omega(i, \boldsymbol{\phi})$ is equal for all i to the overall marginal likelihood of the Markov-APT. The following lemma provides a bottom-up recursive recipe for computing $\xi_A(i, \boldsymbol{\phi})$.

Lemma 2 (Forward-summation)**.** *For $A \in \mathcal{A}^{(\infty)} \backslash \{\Omega\}$,*

$$\xi_A(i, \boldsymbol{\phi}) = \begin{cases} \sum_{i'=1}^{I} \gamma_{i,i'}(A) \cdot M_A^{i'}(\boldsymbol{\theta}_0) \cdot \xi_{A_l}(i', \boldsymbol{\phi}) \xi_{A_r}(i', \boldsymbol{\phi}) & \text{if } n(A) > 1 \\ q_0(\boldsymbol{x}|A) & \text{if } n(A) = 1 \text{ or } A \text{ has no children} \\ 1 & \text{if } n(A) = 0 \end{cases}$$

where $q_0(\boldsymbol{x}|A) = \prod_{x_i \in A} q_0(x_i)/Q_0(A)$ with $q_0 = dQ_0/d\mu$. For $A = \Omega$, we simply replace $\gamma_{i,i'}(A)$ with $\gamma_{i'}(\Omega)$ in the above equation. In particular, $\xi_\Omega(1, \boldsymbol{\phi})$ is the overall marginal likelihood, as a function of the hyperparameters $\boldsymbol{\phi}$.

Remark I: A node A may have no children if Ω is discrete and A is an atom.

Remark II: This lemma shows that one can compute the mapping for A based on those for

its children, A_l and A_r (hence bottom-up). Moreover, one can start the recursion from those As such that $n(A) \leq 1$ but $n(A_p) \geq 2$, because all descendants of such As have no more than one data point and so the mapping is known there.

After computing $\{\xi_A(i, \boldsymbol{\phi}) : A \in \mathcal{A}^{(\infty)} \text{ and } i = 1, 2, \ldots, I\}$, we can then carry out a backward (top-down) recursion to derive the marginal posterior of $\mathcal{C}$.

Theorem 7 (Backward-sampling)**.** *The marginal posterior of the shrinkage states is*

$$\mathcal{C} \,|\, \boldsymbol{\phi}, \boldsymbol{x} \sim \mathrm{MT}(\tilde{\boldsymbol{\gamma}})$$

whose initial state and transition probabilities $\tilde{\boldsymbol{\gamma}} = \{\tilde{\boldsymbol{\gamma}}(A) : A \in \mathcal{A}^{(\infty)}\}$ *are as follows.*

- *The initial state probability vector:*

$$\tilde{\boldsymbol{\gamma}}(\Omega) = \boldsymbol{\gamma}(\Omega)\boldsymbol{D}''(\Omega)/\xi_\Omega(1, \boldsymbol{\phi})$$

- *The state transition probability matrix:*

$$\tilde{\boldsymbol{\gamma}}(A) = \boldsymbol{D}'(A)^{-1}\boldsymbol{\gamma}(A)\boldsymbol{D}''(A)$$

  for all $A \in \mathcal{A}^{(\infty)} \backslash \{\Omega\}$,

where for all $A \in \mathcal{A}^{(\infty)}$, $\boldsymbol{D}'(A)$ *is the* $I \times I$ *diagonal matrix with the diagonal elements being* $\xi_A(i, \boldsymbol{\phi})$ *for* $i = 1, 2, \ldots, I$, *and* $\boldsymbol{D}''(A)$ *is the* $I \times I$ *diagonal matrix with the diagonal elements being* $M_A^i(\boldsymbol{\theta}_0)\xi_{A_l}(i, \boldsymbol{\phi})\xi_{A_r}(i, \boldsymbol{\phi})$ *if* A *has children and* $M_A^i(\boldsymbol{\theta}_0)$ *if not for* $i = 1, 2, \ldots, I$.

Remark: In particular, for any A with $n(A) \leq 1$, by the theorem $\tilde{\boldsymbol{\gamma}}(A) = \boldsymbol{\gamma}(A)$. So *a posteriori* the MT on A with no more than one observation is the same as the prior MT.

We have completely described the three components of the full posterior for $(\boldsymbol{\theta}, \boldsymbol{\nu}, \mathcal{C})$. One can draw from the joint posterior by sampling in the order of $\pi(\mathcal{C} \,|\, \boldsymbol{\phi}, \boldsymbol{x})$, $\pi(\boldsymbol{\nu} \,|\, \boldsymbol{\phi}, \mathcal{C}, \boldsymbol{x})$, and $\pi(\boldsymbol{\theta} \,|\, \boldsymbol{\phi}, \boldsymbol{\nu}, \mathcal{C}, \boldsymbol{x})$. The forward-backward recursion given in Lemma 2 and Theorem 7 is the most computationally intense step in the posterior inference. Fortunately, for any given

data set and prior specification, this recursion only needs to be carried out *once and for all*, because $\tilde{\boldsymbol{\gamma}}$ stays the same for all the posterior draws.

After drawing K posterior samples $(\boldsymbol{\theta}^{(1)}, \boldsymbol{\nu}^{(1)}, \mathcal{C}^{(1)}), (\boldsymbol{\theta}^{(2)}, \boldsymbol{\nu}^{(2)}, \mathcal{C}^{(2)}), \ldots, (\boldsymbol{\theta}^{(K)}, \boldsymbol{\nu}^{(K)}, \mathcal{C}^{(K)})$, one can carry out Bayesian inference in the usual manner. In particular, when one is interested in the unknown distribution Q, one can use $\boldsymbol{\theta}^{(1)}, \boldsymbol{\theta}^{(2)}, \ldots, \boldsymbol{\theta}^{(K)}$ to obtain a posterior sample $Q^{(1)}, Q^{(2)}, \ldots, Q^{(K)}$ while discarding the other variables.

In fact, some posterior summaries can be evaluated analytically without resorting to posterior sampling at all. In particular, the PPD at any $x^* \in \Omega$ is equal to $\xi^*_\Omega(1, \boldsymbol{\phi})/\xi_\Omega(1, \boldsymbol{\phi})$, where $\xi^*_\Omega(1, \boldsymbol{\phi})$ is the overall marginal likelihood computed according to Lemma 2 for a data set that contains the original data plus an additional point at x^*. This is particularly useful in applications such as density estimation because it avoids Monte Carlo errors in computing an estimator. In our numerical examples in Section 3, we compute all PPDs this way. Note that after computing the $\xi_A(i, \boldsymbol{\phi})$ mappings for the original data set, the corresponding mappings $\xi^*_\Omega(i, \boldsymbol{\phi})$ for the new data set can be very quickly obtained by updating only the branch of partition tree in which the new data point x^* falls into, because the mapping stays the same on all other nodes.

2.5 Prior specification for Markov-APTs

Next we provide guidelines for specifying an Markov-APT, i.e. how to choose the hyperparameters $\{(\theta_0(A), \boldsymbol{\gamma}(A), \boldsymbol{F}_{\nu(A)}) : A \in \mathcal{A}^{(\infty)}\}$, in the context of density estimation. The objective is to balance the dual-goal of robustness for a variety of distributional features, and parsimony—involving only a small number of hyperparameters to be set (i.e. the tuning parameters). We give an empirical Bayes strategy to set the tuning parameters adaptively.

Prior choice of $\theta_0(A)$. $\boldsymbol{\theta}_0$ is the PACs corresponding to Q_0, the prior mean of the model. That is, $\theta_0(A) = Q_0(A_l)/Q_0(A)$ for all $A \in \mathcal{A}^{(\infty)}$. Depending on the application, one may or may not have relevant prior knowledge regarding the prior mean. In lack of such knowledge,

a simple default choice is to let Q_0 be uniform over a wide enough interval. Another common situation is that one wants to center the Markov-APT around some parametric family such as the Gaussian location-scale family as in Berger and Guglielmi (2001); Hanson (2006), without specifying exactly which member of that family Q_0 is. This can be achieved by placing another layer of hierarchical prior on Q_0, e.g. on the location and scale parameters (Hanson, 2006), forming a mixture of Markov-APTs.

Prior choice of $\boldsymbol{\gamma}(A)$. In density estimation, we shall choose $\boldsymbol{\gamma}(A)$ to be upper-triangular to achieve adaptive smoothing through stochastically increasing shrinkage. The most simple choice for the transition probabilities that satisfies this condition is

$$\gamma_{i,i'}(A) = \begin{cases} 1/(I-i) & \text{if } i \leq i' \\ 0 & \text{if } i > i' \end{cases}$$

for all $i, i' \in \{1, 2, \ldots, I\}$ and $A \in \mathcal{A}^{(\infty)} \backslash \{\Omega\}$. In other words, given A's parent is in shrinkage state c, then A can take any higher shrinkage state (including c) with equal probability.

This "uniform" transition probability specification is a special case of a more flexible kernel specification. Specifically, we can choose a kernel function $k(i, i')$ such that $k(i, i')$ is non-increasing in $|i - i'|$, and set the transition probability

$$\gamma_{i,i'}(A) \propto k(i, i')\mathbf{1}_{i \leq i'}.$$

A kernel $k(i, i')$ strictly decreasing in $|i - i'|$ will introduce "stickiness" into the shrinkage levels between a node and its parent (and thus also with its siblings and other relatives to a lesser degree). It encourages the shrinkage level to change gradually among nearby nodes in the partition tree, which is particularly useful when the smoothness (or lack of smoothness) of the underlying distribution tends to be similar for places closeby in the sample space.

Of course, typically one does not know *a priori* whether and to what extent such sticky shrinakge is needed. Thus it is useful to allow the stickiness to be adaptively determined.

One natural way to achieving this additional adaptivity is to choose a kernel that contains a tuning parameter for the stickiness, and use empirical Bayes to choose its value. For example, consider the exponential kernel

$$k(i, i') = e^{-\beta|i-i'|},$$

where $\beta \geq 0$ is the stickiness parameter. Note that $\beta = 0$ corresponds to the uniform transition probabilities described above, while a large positive value of β corresponds to strong stickiness in shrinkage. Finally, the initial state probabilities can be simply set to

$$\gamma_1(\Omega) = \gamma_2(\Omega) = \cdots = \gamma_I(\Omega) = \frac{1}{I}.$$

Prior choice of $\boldsymbol{F_\nu}$. Following a common practice (Gelman et al., 2013, Sec. 5.3) in Bayesian hierarchical modeling for Beta-binomial experiments, we specify prior on $\nu(A)$ on the log scale. First, we determine a *global* support for $\log_{10} \nu(A)$, i.e. the union of the supports of all $F^i_{\nu(A)}$. A convenient choice of the global support, aside from the complete shrinkage state included to ensure absolute continuity, is a finite interval $[L, U]$.

A simple and robust strategy for choosing the interval that works in a wide variety of situations is to choose a wide enough range that covers all reasonable shrinkage levels and yet not so wide as to induce excessive prior probability in extremely strong or weak shrinkage levels. We recommend setting $[L, U] = [-1, 4]$. On one end, $\log_{10} \nu(A) = -1$ corresponds to a prior sample size of $10^{-1} = 0.1$, enforcing little shrinkage, while on the other $\log_{10} \nu(A) = 4$ corresponds to shrinkage equivalent to about 10,000 prior "observations" for the local binomial experiment, resulting in very strong shrinkage. We have experimented with treating L and U as tuning parameters and choosing their values in a data dependent fashion using empirical Bayes (described below), but that resulted in little improvement over the very simple choice of [-1,4] in all of the numerical scenarios we investigated.

Given the global support of $\log_{10} \nu(A)$, $[L, U] \cup \{\infty\}$, where ∞ is included for the complete shrinkage state, we now divide this support into I non-overlapping intervals. Specifically, we

let the first $(I-1)$ intervals evenly divide $[L, U]$ and the last being $\{\infty\}$. That is, we have

$$[a(1), a(2)), \;\; [a(2), a(3)), \;\; \ldots, \;\; [a(I-1), a(I)), \;\; \{\infty\}.$$

Then we let

$$F^{i}_{\nu(A)}: \;\; \log_{10} \nu(A) \sim \text{Uniform}(a(i), a(i+1))$$

for $i = 1, 2, \ldots, I-1$ and $F^{I}_{\nu(A)}$ being a point mass at ∞.

Choosing the tuning hyperparameters by empirical Bayes. Our default prior specification is parsimonious in that it reduces the number of free parameters down to two—the number of shrinkage states I and the stickiness parameter β for the transition kernel. One can set these two tuning parameters in a data-adaptive manner by empirical Bayes. In particular, Lemma 2 provides the recipe for computing $\xi_\Omega(1, \boldsymbol{\phi})$, the overall marginal likelihood. We can then compute the marginal likelihood as a function of the tuning parameters $\xi_\Omega(1, \boldsymbol{\phi}(I, \beta))$. Maximizing this likelihood over a grid of allowed values produces the maximum marginal likelihood estimate (MMLE) $(\hat{I}, \hat{\beta})$, which one can then keep fixed in the inference.

3 Performance evaluation in density estimation

In this section we evaluate the performance of the Markov-APT in density estimation under five schematic simulation scenarios (Figure 4). Each scenario corresponds to an underlying density with a particular type of structure commonly encountered in real applications. For each scenario, we simulate data sets of six different sample sizes—125, 250, 500, 750, 1000, and 1250. We compare the performance of Markov-APT to that of three other nonparametric models—namely, the PT, the OPT, and the DPM of normals (Escobar and West, 1995). We fit the DPM in `R` using the library `DPpackage` (Jara, 2007; Jara et al., 2011). Details on the specification of the DPM are given in Supplementary Materials S2.

For each method, we use the PPD, denoted by $\hat{f}$, as an estimator. To measure performance, we adopt the L_1 loss, i.e. the L_1 distance between $\hat{f}$ and the true density f_0,

$||\hat{f} - f_0||_1 = \int |\hat{f} - f_0| d\mu$. For each simulation scenario and sample size, we generate $K = 200$ data sets, with $\hat{f}^{(k)}$ being the estimate for the kth data set. We numerically calculate $||\hat{f}^{(k)} - f_0||_1$ using Riemann integral for all four methods. To make a comparison, for each of the competitors—PT, OPT, and DPM—we compute the percentage difference between its L_1 loss and that of the Markov-APT:

$$\frac{||\hat{f}^{(k)}_{\text{Competitor}} - f_0||_1 - ||\hat{f}^{(k)}_{\text{Markov-APT}} - f_0||_1}{||\hat{f}^{(k)}_{\text{Markov-APT}} - f_0||_1} \times 100\%.$$

A positive percentage increase indicates an outperformance of the Markov-APT over the competitor, with larger values indicating more effectiveness of the Markov-APT. Computing this measure of relative performance for each simulation allows us to evaluate both the average performance increase and the variability in the improvement across repeated experiments. In addition, we also estimate the average L_1 loss, i.e. the L_1 risk, $R_{L_1}(f_0, \hat{f}) = E_{f_0}||\hat{f} - f_0||_1$ for each method under each simulation setting using the Monte Carlo average

$$\widehat{R}_{L_1}(f_0, \hat{f}) = \frac{1}{K} \sum_{k=1}^{K} ||\hat{f}^{(k)} - f_0||_1.$$

In all of the simulation settings, we adopt the prior specification recommended in Section 2.5 with an exponential transition kernel, and use empirical Bayes to set the tuning parameters (I, β). The range of tuning parameter values over which we maximize the marginal likelihood is $\{2, 3, \ldots, 11\} \times [0, 2]$. The OPT also involves a tuning parameter ρ_0, the prior "stopping" probability (Wong and Ma, 2010), which we set by empirical Bayes using MMLE over $[0, 1]$. We implement the Markov-APT, PT, and OPT models up to the 12th level in the partition tree. Deeper partition trees result in little numerical difference.

Figure 4 presents the true densities for all scenarios from which the data are simulated. In each scenario, the underlying density is supported on the interval [0,1]. Our proposed model does not require the support to be a bounded interval, and this choice here is to simplify the numerical evaluation of the L_1 loss. This causes no loss of generality in our

simulation because any density on $\mathbb{R}$ can be transformed onto [0,1] after applying, say, a cdf transformation. In Figure 5 we present histograms of the percentage increase in L_1 loss for the three competitor methods relative to Markov-APT for each simulation setting. For easier comparison, we overlay the histograms for three different sample sizes—125 (small), 500 (medium), and 1000 (large)—to show how the relative performance changes for different sample sizes. (We have chosen to only show the three sample sizes in this figure rather than all six sample sizes because overlaying six histograms makes the plot illegible while using the three sample sizes is sufficient to convey the main finding.) Figure 6 presents the L_1 risks for all methods and scenarios versus sample size. Finally, to help understand why each method performs well or poorly in each scenario, in Figure 7 we plot a typical PPD for each model under each scenario for a sample size that well differentiates the performance of the methods.

- *Scenario 1: Spiky local structures.* In this case, the true distribution is

$$0.2\,\mathrm{U}(0,1) + 0.2\,\mathrm{U}(0.2,0.205) + 0.2\,\mathrm{U}(0.4,0.405) + 0.2\,\mathrm{U}(0.6,0.605) + 0.2\,\mathrm{U}(0.8,0.805).$$

See Figure 4(a) for the true density. This represents the case when the underlying distribution has a few spiky structures in the midst of a flat background. In this case the key is to effectively determine the size (or height) of those spikes and pin down their boundaries.

Multi-resolution inference methods such as wavelets, PT, OPT, and our Markov-APT are particularly effective in capturing abrupt changes such as spikes and sharp boundaries in the nonparametric quantity of interest. Thus scenarios where the underlying distribution predominantly consists of spiky structures are the most favorable scenario for such methods in comparison to mixture-based method such as the DPM.

Indeed, as shown in Figure 5(a) and Figure 6(a), the Markov-APT performs substantially better than DPM, especially for medium and large sample sizes. It may first appear surprising that the PT, being a multi-resolution approach, performs the worst among all, in fact substantially worse than DPM. But this can be expected because the PT is unable to

amply capture the height of the spikes due to over shrinkage at high-resolutions. In contrast, the amount of shrinkage under the Markov-APT is adaptive and thus automatically adjusts to low levels in and around the spikes. Interestingly, the OPT, which only allows no shrinkage or complete shrinkage, performs even slightly better than the Markov-APT. But in fact this is not surprising at all. When the underlying density is a step function as in the current scenario, the only appropriate shrinkage levels are indeed no shrinkage and complete shrinkage. Therefore the OPT is in fact an "oracle" in this case and one should expect it to perform the best. It is reassuring to see that the Markov-APT, while allowing much more flexible shrinkage levels, did not lose much efficiency relative to the "oracle".

From the PPDs in Figure 7(a) we see that the DPM tends to overestimate the height of the spikes—this is likely because in order to the characterize the sharp boundaries the DPM needs to "squeeze" the mixture component to have very thin tails, and thus making the mode of the mixture component much taller than the truth.

- *Scenario 2: Non-overlapping structures of different scales.* The true distribution is

$$0.1\,\mathrm{U}(0,1) + 0.3\,\mathrm{U}(0.25,0.5) + 0.4\,\mathrm{Beta}_{(0.25,0.5)}(2,2) + 0.2\,\mathrm{Beta}(6000,4000)$$

See Figure 4(b) for the true density. This is the scenario given earlier in Example 1. The underlying density has two bumps of different scales—one large and the other small—that are not overlapping with each other. The high-resolution, spiky structure now has a smooth boundary in contrast to the abrupt boundaries in the previous scenario. This is a favorable scenario for kernel mixture methods such as the DPM. By allowing the local kernel to have varying variance, the DPM is also able to adapt to the different scales of the two bumps. Thus one would expect the DPM to perform well.

From Figure 5(b) and Figure 6(b), we see that the Markov-APT achieves better performance than the DPM at sample size 125, and comparable performance at larger sample sizes. The performance gain of Markov-APT over the OPT and PT is substantial at all sam-

ple sizes. Figure 7(b) shows that the OPT substantially oversmooths the large-scale feature while capturing the small-scale feature well, and the PT oversmooths the small-scale feature while undersmooths at high-resolutions within the large-scale feature.

- *Scenario 3: Overlapping structures of different scales.* The true distribution is

$$0.1\,\mathrm{U}(0,1) + 0.3\,\mathrm{U}(0.25,0.5) + 0.4\,\mathrm{Beta}_{(0.25,0.5)}(2,2) + 0.2\,\mathrm{Beta}(4000,6000).$$

See Figure 4(c) for the true density. This case is similar to the previous except that now the spiky local structure lies inside the smooth large-scale structure. From Figure 6(c), we see that the performance of the Markov-APT is essentially unchanged from the case where the structures are non-overlapping. In contrast, Figure 5(c) and Figure 6(c) show that the other methods all perform quite differently in this scenario. In particular, there is a substantial decay in performance for the DPM at smaller sample sizes compared to the case with non-overlapping structures, whereas the OPT and PT perform better for the current scenario. The dramatic change in the performance of PT compared to that in Scenario 2 illustrates the importance of adaptivity in achieving robust inference. Non-adaptive methods may be performing well in some situations but very poorly in another with only small or modest changes in the underlying distribution.

From Figure 7(c) we see that the three multi-resolution methods are all capable of characterizing the jump in the density even with small sample sizes. The sudden improvement in PT's performance is readily explained in Figure 7(c). The estimation error from PT in this and the previous scenario comes from two sources—the under-smoothing (i.e. under-shrinkage) in the large-scale smooth structure and over-smoothing (i.e. over-shrinkage) in the local spiky structure. By moving the spiky structure into the smooth structure, the error that comes from the under-smoothing in the large-scale structure is reduced because the smooth portion of the large-scale structure now accounts for a smaller proportion of the total probability mass, while the error that comes from over-smoothing the local structure

is also reduced because now the local structure corresponds to more probability mass and hence more data. Note also in Figure 5(c) and Figure 6(c) that the relative performance gain of Markov-APT over PT increases with the sample size.

- *Scenario 4: Sharp boundaries.* The true distribution is

$$0.1\text{Beta}(2,2)+0.25\text{U}(0.3,0.55)+0.05\text{Beta}_{(0.3,0.55)}(2,2)+0.55\text{U}(0.55,0.8)+0.05\text{Beta}(0.55,0.8).$$

See Figure 4(d) for the true density. In this scenario the underlying density is a couple of smooth structures with sharp boundaries separating them. In order to characterize the sharp jumps, the DPM introduces several mixture components, resulting in the loss of performance. Figure 5(d) and Figure 6(d) show that the Markov-APT and the OPT perform the best and in very similar way over the entire range of sample sizes. From Figure 7(d) we see that the Markov-APT is able to both capture the sharp boundaries and the smooth modes, while the OPT again tends to oversmooth the two local modes, but is able to capture the sharp boundaries accurately. The L_1 loss in this example is predominantly contributed from the boundaries, and so the oversmoothing does not impair much of the performance of the OPT. On the other hand, the PT again performs significantly worse than any other method due to its substantial undersmoothing. At small sample sizes, the DPM achieves comparable performance as the Markov-APT, but the relative performance gain of the Markov-APT over the DPM is widened for larger sample sizes. From Figure 7(d) we see that in order to characterize the sharp boundaries, the DPM needs to introduce a number of additional mixture components, resulting in overfitting in regions away from the boundaries.

- *Scenario 5: Globally smooth structure.* The true distribution is $\text{Beta}(10,20)$.

See Figure 4(e) for the true density. In this case the underlying density is an approximately Gaussian smooth distribution, and DPM with a Gaussian kernel is essentially the true model for this scenario, and unsurprisingly performs the best. Such globally smooth distributions are the "least favorable" scenario for the three PT-type multi-resolution methods. Figure 5(e)

and Figure 6(e) show that for all sample sizes, the L_1 loss is on average about 50% smaller under the DPM vs the Markov-APT. Among the multi-resolution methods, the Markov-APT substantially outperforms the OPT and the PT, and the performance gain widens for larger sample sizes.

4 Discussion

We have showed that inference under PT-type multi-resolution models can be understood from a shrinkage perspective, and have introduced a hierarchical Bayesian approach to incorporating adaptive shrinkage. The APT and Markov-APT models can be easily applied in hypothesis testing such as for testing a parametric null versus a nonparametric alternative as previously studied in Berger and Guglielmi (2001) and testing two-sample differences as studied in Holmes et al. (2009); Ma and Wong (2011); Chen and Hanson (2014). Under this framework, the testing of features of the underlying distribution is transformed into testing a collection of local hypotheses organized on a partition tree, one for each node in the tree. As such, stochastically increasing shrinkage should be of less importance (as a motivation for choosing the appropriate probability transition matrix of the MT) than proper adjustment of multiple testing. Thus in such applications the strategy for prior specification is different.

Last but very importantly, inference under PT-type multi-resolution models such as PT, OPT, APT and Markov-APT is extremely computationally efficient due to the conjugate hierarchical design and the forward-backward algorithm. In particular, for the most sophisticated model in this class, the Markov-APT, the dominating step in computing the posterior and PPD is computing the mappings $\xi_A(i, \boldsymbol{\phi})$ through Lemma 2, but it takes less than 0.1 second even with sample size 1250 for all of our numerical scenarios on a single Intel Core-i7 3.6Ghz CPU core with 400 Mbs of RAM. Moreoever, the computing time and required RAM stay essentially constant for sample sizes in the typical range (from tens to tens of thousands). In contrast, fitting the DPM model using MCMC is much more computation-

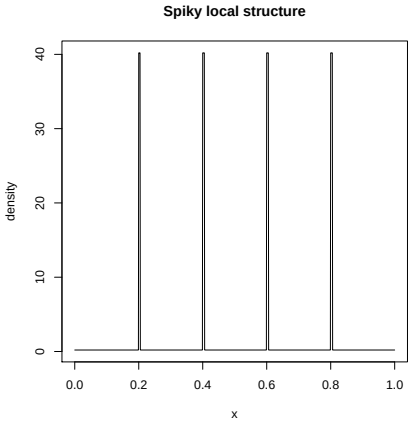

(a) Scenario 1

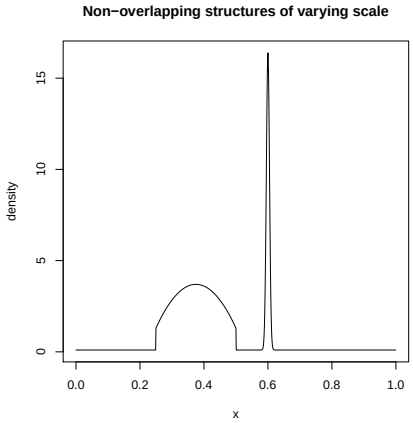

(b) Scenario 2

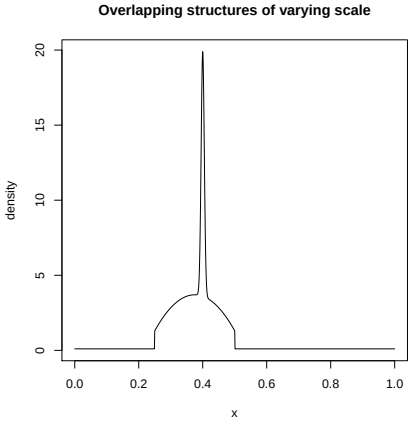

(c) Scenario 3

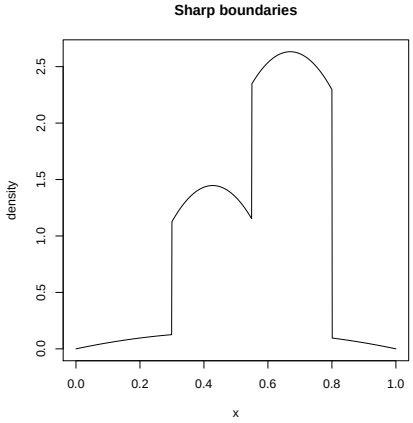

(d) Scenario 4

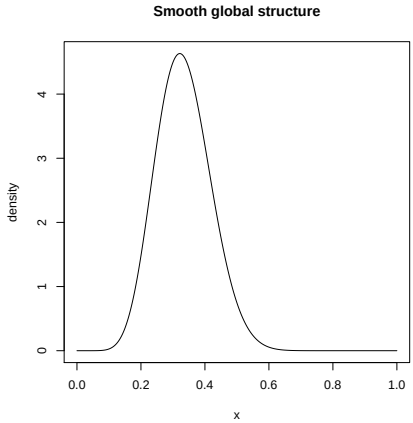

(e) Scenario 5

Figure 4: True densities of the five simulation scenarios

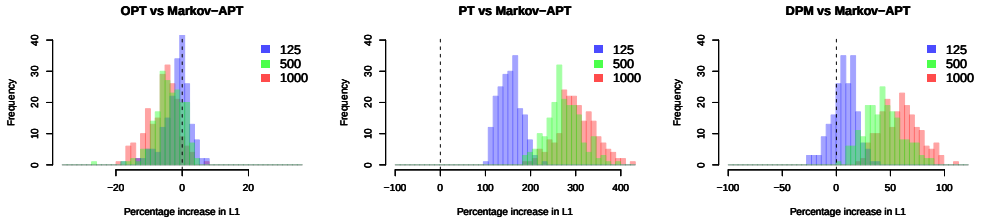

(a) Scenario 1. Spiky local structure.

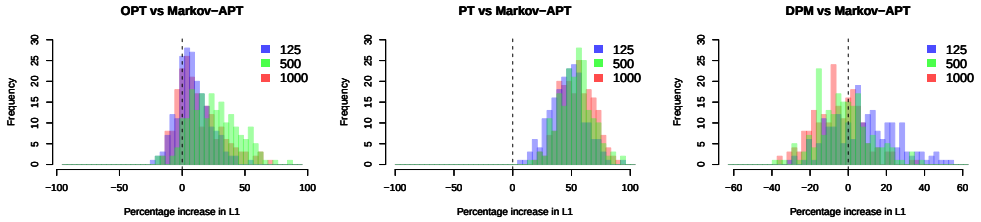

(b) Scenario 2. Non-overlapping structures of varying scale.

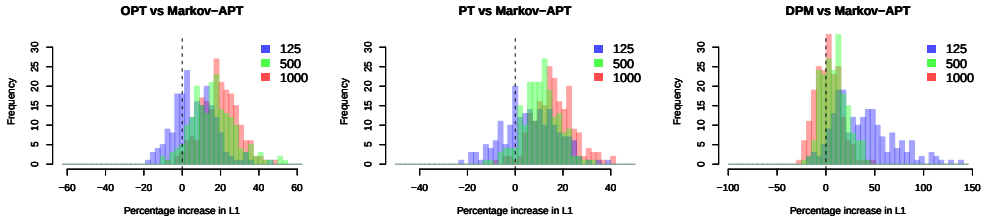

(c) Scenario 3. Overlapping structures of varying scale.

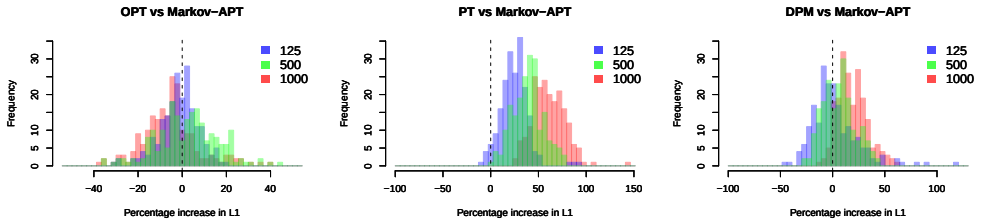

(d) Scenario 4. Sharp boundaries.

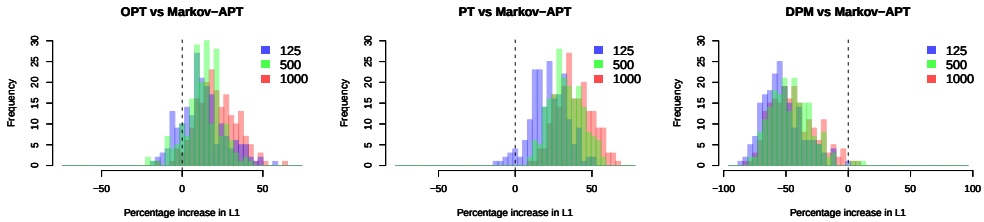

(e) Scenario 5. Smooth global structure.

Figure 5: Histograms of percentage increase in L_1 risk for three methods compared to Markov-APT over three sample sizes—125, 500, and 1000.

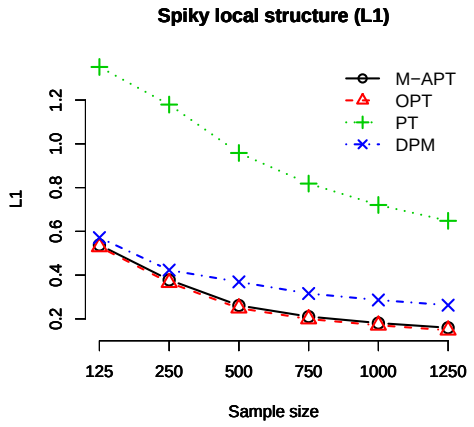

(a) Scenario 1

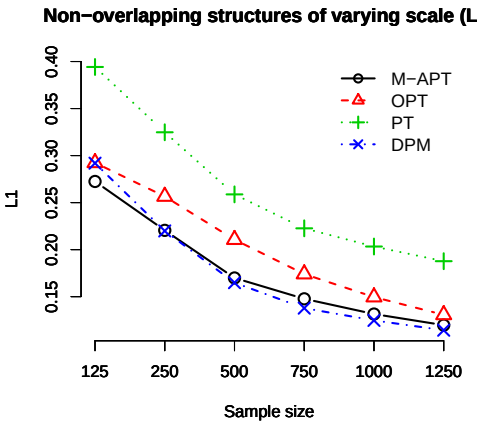

(b) Scenario 2

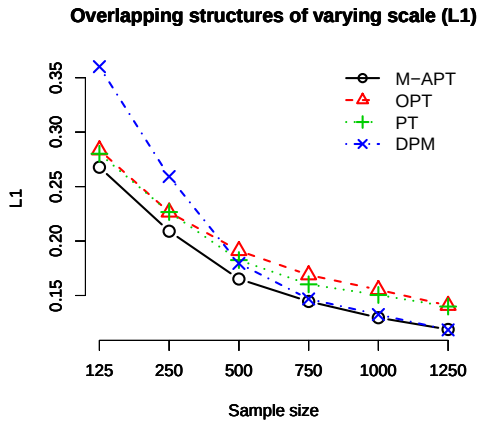

(c) Scenario 3

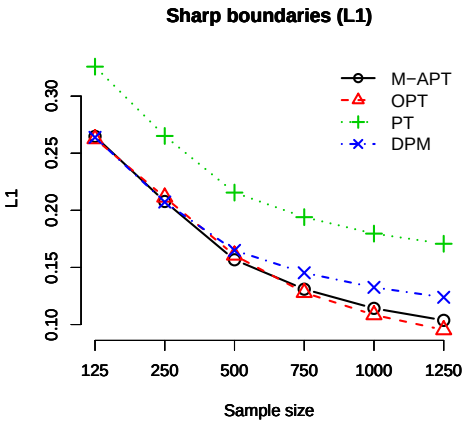

(d) Scenario 4

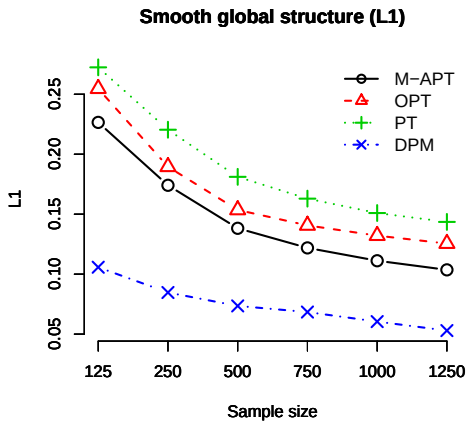

(e) Scenario 5

Figure 6: Estimated L_1 risk of four methods by sample size for the five simulation scenarios.

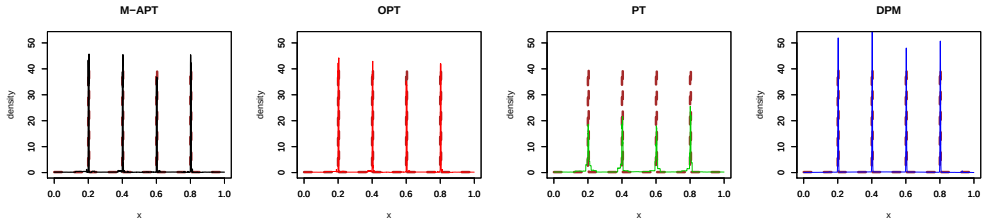

(a) Scenario 1. $n = 500$. Markov-APT: $\hat{I} = 4$ and $\hat{\beta} = 1.2$. OPT: $\hat{\rho} = 0.26$.

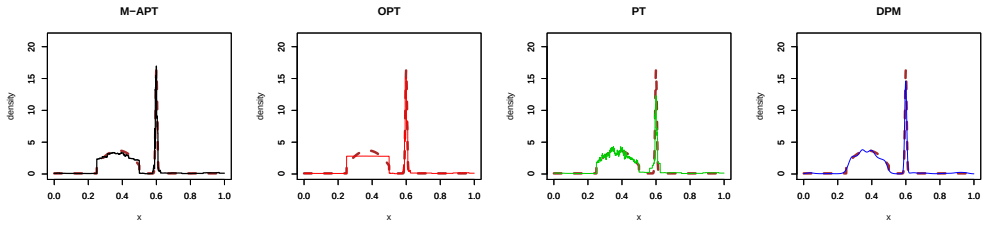

(b) Scenario 2. $n = 500$. Markov-APT: $\hat{I} = 8$ and $\hat{\beta} = 0.70$. OPT: $\hat{\rho} = 0.38$.

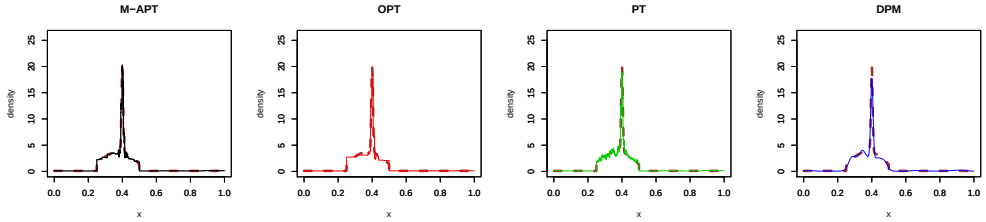

(c) Scenario 3. $n = 500$. Markov-APT: $\hat{I} = 11$ and $\hat{\beta} = 0.50$. OPT: $\hat{\rho} = 0.50$.

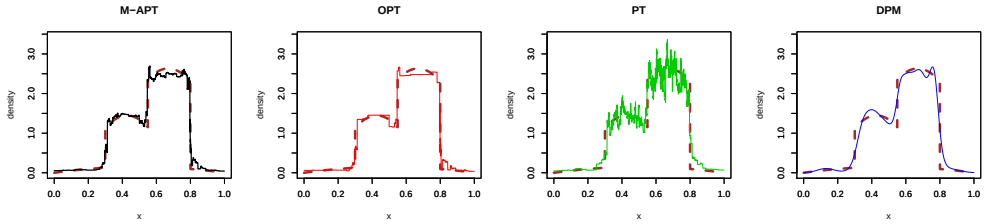

(d) Scenario 4. $n = 1000$. Markov-APT: $\hat{I} = 6$ and $\hat{\beta} = 0.75$. OPT: $\hat{\rho} = 0.42$.

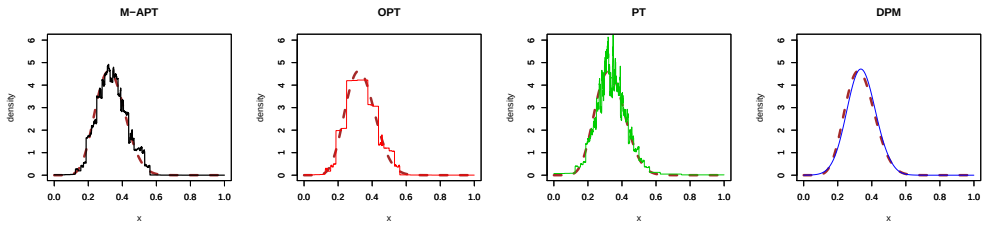

(e) Scenario 5. $n = 500$. Markov-APT: $\hat{I} = 11$ and $\hat{\beta} = 0.60$. OPT: $\hat{\rho} = 0.38$.

Figure 7: Typical PPDs (solid) of four methods and true density (dashed) for the five simulation scenarios. Sample sizes and tuning parameter values chosen by MMLE are given.

ally expensive both in time and in memory. In particular, fitting the DPM in `R` using the `DPpackage` for each of the five scenarios at sample size 1250 takes about 4 to 7 minutes on the same machine and requires about 1.5 Gbs of RAM.

We believe that PT-type multi-resolution methods has tremendous potential for applications where the underlying distribution involves abrupt changes such as spikes or sharp boundaries, as well as where computational efficiency is of critical importance, such as in real-time change-point detection, online applications, and those applications with very large sample sizes. Therefore, additional effort is worthwhile to study the theory and further improve the statistical and computational performance of this class of methods.

Acknowledgment

This research is supported by NSF grant DMS-1309057.

References

Abramovich, F. and Y. Benjamini (1995). Thresholding of wavelet coefficients as multiple hypotheses testing procedure. In A. Antoniadis and G. Oppenheim (Eds.), *Wavelets and Statistics*, Volume 103 of *Lecture Notes in Statistics*, pp. 5–14. Springer New York.

Abramovich, F., T. Sapatinas, and B. W. Silverman (1998). Wavelet thresholding via a bayesian approach. *Journal of the Royal Statistical Society: Series B 60*(4), 725–749.

Antoniadis, A. and T. Sapatinas (2001). Wavelet shrinkage for natural exponential families with quadratic variance functions. *Biometrika 88*(3), 805–820.

Berger, J. O. and A. Guglielmi (2001). Bayesian and conditional frequentist testing of a parametric model versus nonparametric alternatives. *Journal of the American Statistical Association 96*(453), 174–184.

Brown, P. J., T. Fearn, and M. Vannucci (2001). Bayesian wavelet regression on curves with application to a spectroscopic calibration problem. *Journal of the American Statistical Association 96*(454), 398–408.

Cai, T. and L. D. Brown (1999). Wavelet estimation for samples with random uniform design. *Statistics & Probability Letters 42*(3), 313 – 321.

Cai, T. T. and L. D. Brown (1998, 10). Wavelet shrinkage for nonequispaced samples. *The Annals of Statistics 26*(5), 1783–1799.

Chang, S., B. Yu, and M. Vetterli (2000, Sep). Adaptive wavelet thresholding for image denoising and compression. *Image Processing, IEEE Transactions on 9*(9), 1532–1546.

Chen, Y. and T. E. Hanson (2014). Bayesian nonparametric k-sample tests for censored and uncensored data. *Computational Statistics & Data Analysis 71*(0), 335 – 346.

Chipman, H. A., E. D. Kolaczyk, and R. E. Mcculloch (1997). Adaptive Bayesian wavelet shrinkage. *Journal of The American Statistical Association 92*, 1413–1421.

Clyde, M. and E. I. George (2000a). Flexible empirical bayes estimation for wavelets. *Journal of the Royal Statistical Society: Series B 62*(4), 681–698.

Clyde, M. and E. I. George (2000b). Flexible empirical bayes estimation for wavelets. *Journal of the Royal Statistical Society: Series B 62*(4), 681–698.

Clyde, M., G. Parmigiani, and B. Vidakovic (1998). Multiple shrinkage and subset selection in wavelets. *Biometrika 85*(2), 391–401.

Crouse, M. S., R. D. Nowak, and R. G. Baraniuk (1998, April). Wavelet-based statistical signal processing using hidden Markov models. *Signal Processing, IEEE Transactions on 46*(4), 886–902.

Donoho, D. (1995, May). De-noising by soft-thresholding. *Information Theory, IEEE Transactions on 41*(3), 613–627.

Donoho, D. and I. M. Johnstone (1995). Adapting to unknown smoothness via wavelet shrinkage. *Journal of the American Statistical Association 90*, 1200–1224.

Donoho, D. L. and I. M. Johnstone (1994). Ideal spatial adaptation via wavelet shrinkage. *Biometrika* (81), 425–455.

Donoho, D. L. and I. M. Johnstone (1998, 06). Minimax estimation via wavelet shrinkage. *The Annals of Statistics 26*(3), 879–921.

Donoho, D. L., I. M. Johnstone, G. Kerkyacharian, and D. Picard (1995). Wavelet shrinkage: Asymptopia? *Journal of the Royal Statistical Society. Series B 57*(2), pp. 301–369.

Donoho, D. L., I. M. Johnstone, G. Kerkyacharian, and D. Picard (1996). Density estimation by wavelet thresholding. *The Annals of Statistics 24*(2), pp. 508–539.

Draper, D. (1999). Discussion of "Bayesian nonparametric inference for random distributions and related functions". *Journal of the Royal Statistical Society Series B* (61), 510–513.

Escobar, M. D. and M. West (1995). Bayesian density estimation and inference using mixtures. *Journal of the American Statistical Association 90*, 577–588.

Ferguson, T. S. (1973). A Bayesian analysis of some nonparametric problems. *Ann. Statist. 1*, 209–230.

Ferguson, T. S. (1974). Prior distributions on spaces of probability measures. *Ann. Statist. 2*, 615–629.

Gelman, A., J. Carlin, H. Stern, D. Dunson, A. Vehtari, and D. Rubin (2013). *Bayesian Data Analysis, Third Edition*. Chapman & Hall/CRC Texts in Statistical Science. Taylor & Francis.

Ghosh, J. K. and R. V. Ramamoorthi (2003). *Bayesian Nonparametrics*. Springer Series in Statistics. New York: Springer-Verlag.

Hall, P., G. Kerkyacharian, and D. Picard (1998). Block threshold rules for curve estimation using kernel and wavelet methods. *The Annals of Statistics 26*(3), pp. 922–942.

Hanson, T. and W. O. Johnson (2002). Modeling regression error with a mixture of Pólya trees. *Journal of the American Statistical Association 97*(460), pp. 1020–1033.

Hanson, T. E. (2006). Inference for mixtures of finite Pólya tree models. *J. Amer. Statist. Assoc. 101*(476), 1548–1565.

Herrick, D. R. M., G. P. Nason, and B. W. Silverman (2001). Some New Methods for Wavelet Density Estimation. *Sankhya: The Indian Journal of Statistics, Series A 63*(3), 394–411.

Holmes, C. C., F. Caron, J. E. Griffin, and D. A. Stephens (2009, Oct). Two-sample Bayesian nonparametric hypothesis testing.

Jara, A. (2007). Applied Bayesian non- and semi-parametric inference using DPpackage. *R News 7*(3), 17–26.

Jara, A., T. Hanson, F. Quintana, P. Müller, and G. Rosner (2011). DPpackage: Bayesian semi- and nonparametric modeling in R. *Journal of Statistical Software 40*(5), 1–30.

Johnstone, I. M. (1999). Wavelet shrinkage for correlated data and inverse problems: Adaptivity results.

Johnstone, I. M. and B. W. Silverman (1997). Wavelet threshold estimators for data with correlated noise. *Journal of the Royal Statistical Society: Series B 59*(2), 319–351.

Johnstone, I. M. and B. W. Silverman (2005). Empirical bayes selection of wavelet thresholds. *The Annals of Statistics 33*(4), pp. 1700–1752.

Kolaczyk, E. D. (1999). Wavelet shrinkage estimation of certain Poisson intensity signals using corrected thresholds. *Statistica Sinica 9*, 119–135.

Koo, J.-Y. and C. Kooperberg (2000). Logspline density estimation for binned data. *Statistics & Probability Letters 46*(2), 133 – 147.

Kraft, C. H. (1964). A class of distribution function processes which have derivatives. *J. Appl. Probability 1*, 385–388.

Lavine, M. (1992). Some aspects of Pólya tree distributions for statistical modelling. *Ann. Statist. 20*(3), 1222–1235.

Lavine, M. (1994). More aspects of Pólya tree distributions for statistical modelling. *Ann. Statist. 22*(3), 1161–1176.

Liu, J. S. (2001, January). *Monte Carlo Strategies in Scientific Computing*. Springer.

Ma, L. and W. H. Wong (2011). Coupling optional Pólya trees and the two sample problem. *Journal of the American Statistical Association 106*(496), 1553–1565.

Mauldin, D. R., W. D. Sudderth, and S. C. Williams (1992). Pólya trees and random distributions. *Ann. Statist. 20*(3), 1203–1221.

Morris, J. S. and R. J. Carroll (2006). Wavelet-based functional mixed models. *Journal of the Royal Statistical Society: Series B 68*(2), 179–199.

Moulin, P. and J. Liu (1999, Apr). Analysis of multiresolution image denoising schemes using generalized gaussian and complexity priors. *Information Theory, IEEE Transactions on 45*(3), 909–919.

Nason, G. P. (1996). Wavelet shrinkage using cross-validation. *Journal of the Royal Statistical Society. Series B 58*(2), pp. 463–479.

Nason, G. P. and B. W. Silverman (1995). The stationary wavelet transform and some statistical applications. pp. 281–300. Springer-Verlag.

Portilla, J., V. Strela, M. Wainwright, and E. Simoncelli (2003, Nov). Image denoising using scale mixtures of gaussians in the wavelet domain. *Image Processing, IEEE Transactions on 12*(11), 1338–1351.

Romberg, J., H. Choi, and R. Baraniuk (2001, Jul). Bayesian tree-structured image modeling using wavelet-domain hidden markov models. *Image Processing, IEEE Transactions on 10*(7), 1056–1068.

Vannucci, M. (1995). Nonparametric density estimation using wavelets. Discussion paper 95-26, Duke University.

Vannucci, M. and F. Corradi (1999). Covariance structure of wavelet coefficients: Theory and models in a bayesian perspective. *Journal of the Royal Statistical Society. Series B 61*(4), pp. 971–986.

Vidakovic, B. (1998). Nonlinear wavelet shrinkage with bayes rules and bayes factors. *Journal of the American Statistical Association 93*, 173–179.

Vidakovic, B. (1999). *Statistical Modeling by Wavelets* (1st ed.). Hoboken, NJ, USA: John Wiley & Sons, Inc.

Walker, S. G., P. Damien, P. W. Laud, and A. F. M. Smith (1999). Bayesian nonparametric inference for random distributions and related functions. *Journal of the Royal Statistical Society: Series B 61*(3), 485–527.

Wong, W. H. and L. Ma (2010). Optional Pólya tree and Bayesian inference. *Annals of Statistics 38*(3), 1433–1459.

Supplementary Materials

S1. Technical proofs

Proof of Lemma 1. The existence of a collection of PACs for G follows immediately from the definition of PACs by letting $\theta(A) = G(A_l)/G(A)$ for each A such that $G(A) > 0$ and $\theta(A) = 0$ otherwise. The uniqueness follows because $\mathcal{A}^{(\infty)}$ forms a π-system that generates the Borel σ-algebra. So by the extension theorem, two distributions with the same PACs on all As such that $G(A) > 0$ must be the same up to a set of μ-measure 0. □

Proof of Theorem 1. Given $\boldsymbol{\nu}$, Q has a PT distribution with mean Q_0. That is $E(Q(B)|\boldsymbol{\nu}) = Q_0(B)$. The result follows immediately by the law of iterated expectation. □

Proof of Theorem 2. This theorem follows immediately from applying the Bayes rule to the conjugate hierarchical model, and is described in detail in Section 2.4. □

Proof of Theorem 3. Given $\mathcal{C}$, Q has an APT distribution with mean Q_0 by Theorem 1. That is $E(Q(B)|\mathcal{C}) = Q_0(B)$. The result follows immediately by the law of iterated expectation. □

Proof of Theorem 4. Let $Q^{(k)}$ be the level-k truncated version of Q. That is, $Q^{(k)}(A) = Q(A)$ for all $A \in \mathcal{A}^{(k)}$ and $Q^{(k)}(\cdot|A) = Q_0(\cdot|A)$ for all $A \in \mathcal{A}^k$. By the same argument as in the proof of Theorem 1 in Wong and Ma (2010) (with Q_0 replacing μ), we know that $Q^{(k)}$ converges in total variational distance to Q as $k \to \infty$. By construction, $Q^{(k)} \ll Q_0$ for all k. Now for any set B such that $Q(B) > 0$, then there must exist some k such that $Q^{(k)}(B) > 0$, and therefore $Q_0(B) > 0$. Hence $Q \ll Q_0$. □

Proof of Theorem 5. Let $\tilde{q} = q/q_0$ and $\tilde{g} = g/q_0$ where $q_0 = dQ_0/d\mu$. Our goal is to prove that for any $\tau > 0$,

$$P\left(\int |\tilde{q} - \tilde{g}| dQ_0 < \tau\right) > 0.$$

First we assume that $\tilde{g}$ is continuous and bounded, and let M be a finite upperbound of $\tilde{g}$. For any $\sigma > 0$, there exists a compact set E such that there is a partition $\Omega = \cup_i A_i$ such that the diameter of each $A_i \cap E$ is less than σ. By the absolute continuity of G w.r.t Q_0, there exists $\beta(\sigma) > 0$ such that $G(E^c) < \beta(\sigma)$ if $Q_0(E^c) < \sigma$ and $\beta(\sigma) \downarrow 0$ as $\sigma \downarrow 0$. We define the modulus of continuity of $\tilde{g}$ on E as

$$\delta_E(\epsilon) = \sup_{x,y \in E: |x-y|<\epsilon} |\tilde{g}(x) - \tilde{g}(y)|.$$

Note that by the continuity of $\tilde{g}$ and the compactness of E, $\delta_E(\epsilon) \downarrow 0$ as $\epsilon \downarrow 0$. Now we approximate $\tilde{g}$ by a step function $\tilde{g}^*(x) = \sum_i \tilde{g}_i^* I_{A_i}$ where $\tilde{g}_i^* = \int_{A_i \cap E} \tilde{g} dQ_0 / Q_0(A_i \cap E)$. Let $D_\epsilon(\tilde{g})$ be the set of step functions $h(\cdot) = \sum_i h_i I_{A_i}(\cdot)$ such that $\sup_i |h_i - \tilde{g}_i^*| < \delta_E(\epsilon) + M\sigma$.

Suppose $h \in D_\epsilon(\tilde{g})$. For any $B \in \mathcal{B}$, the Borel sets, we have $B_i = B \cap A_i$. Then

$$\begin{aligned}
&\left| \int_B (h - \tilde{g}) dQ_0 \right| \leq \sum_i |h_i - \tilde{g}_i^*| Q_0(B_i) + \sum_i \left| \tilde{g}_i^* \, Q_0(B_i) - \int_{B_i} \tilde{g} dQ_0 \right| \\
\leq &(\delta_E(\epsilon) + M\sigma) Q_0(B) + \sum_i \left| \tilde{g}_i^* \, Q_0(B_i \cap E) - \int_{B_i \cap E} \tilde{g} dQ_0 \right| + \sum_i \left| \tilde{g}_i^* \, Q_0(B_i \cap E^c) - \int_{B_i \cap E^c} \tilde{g} dQ_0 \right| \\
\leq &(\delta_E(\epsilon) + M\sigma) Q_0(B) + \sum_i r_i + 2M \cdot Q_0(E^c) \\
< &(\delta_E(\epsilon) + M\sigma) Q_0(B) + \sum_i r_i + 2M\sigma
\end{aligned}$$

where

$$\begin{aligned}
r_i &= Q_0(B_i \cap E) \left| \frac{\int_{A_i \cap E} \tilde{g} dQ_0}{Q_0(A_i \cap E)} - \frac{\int_{B_i \cap E} \tilde{g} dQ_0}{Q_0(B_i \cap E)} \right| \\
&= Q_0(B_i \cap E) \left| \frac{\int_{A_i \cap E} (\tilde{g}(x) - \tilde{g}(x_i)) \, q_0(x) \, dx}{Q_0(A_i \cap E)} - \frac{\int_{B_i \cap E} (\tilde{g}(x) - \tilde{g}(x_i)) \, q_0(x) \, dx}{Q_0(B_i \cap E)} \right|
\end{aligned}$$

for some $x_i \in B_i$. Thus

$$|r_i| < 2\delta_E(\epsilon) Q_0(B_i)$$

and so

$$\Big|\int_B (h-\tilde{g})dQ_0\Big| < 3\delta_E(\epsilon)Q_0(B) + 3M\sigma \quad \text{for all } B \in \mathcal{B}.$$

Therefore by taking $B = \{x : h > \tilde{g}\}$ and $B = \{x : h \leq \tilde{g}\}$, we get

$$\int |h-\tilde{g}|dQ_0 < 3\delta_E(\epsilon) + 6M\sigma.$$

Now we let $Q^{(k)}$ be the level-k truncated version of Q. That is, $Q^{(k)}(A) = Q(A)$ for all $A \in \mathcal{A}^{(k)}$ and $Q^{(k)}(\cdot|A) = Q_0(\cdot|A)$ for all $A \in \mathcal{A}^k$. By the conditions in the theorem, we have for $\tilde{q}^{(k)} = q^{(k)}/q_0$ where $q^{(k)} = dQ^{(k)}/d\mu$,

$$P\left(\tilde{q}^{(k)} \in D_\epsilon(\tilde{g}) \text{ for all large } k\right) > 0.$$

Thus

$$P\left(\int |\tilde{q}^{(k)} - \tilde{g}|dQ_0 < 3\delta_E(\epsilon) + 6M\sigma \text{ for all large } k\right) > 0.$$

But since

$$P\left(\int |\tilde{q}^{(k)} - \tilde{q}|dQ_0 \to 0\right) = 1,$$

combining these we get

$$P\left(\int |\tilde{q} - \tilde{g}|dQ_0 < 4\delta_E(\epsilon) + 6M\sigma\right) > 0.$$

The result follows by letting $\epsilon \downarrow 0$ and $\sigma \downarrow 0$.

Finally, if $\tilde{g}$ is not continuous and bounded, then since Q_0 is a probability measure, $\tilde{g}$ can be approximately arbitrarily well in L_1 w.r.t Q_0 by a continuous bounded density. □

Proof of Theorem 6. Let $p_0 = dP_0/d\mu$, $q_0 = dQ_0/d\mu$, $\tilde{p}_0 = dP_0/dQ_0$, and for any $Q \ll Q_0$, $\tilde{q} = dQ/dQ_0$. Let M be a finite upperbound of $\tilde{p}_0$. Then the Kullback-Leibler (K-L) distance

between p_0 and q is given by

$$\mathrm{KL}_\mu(p_0, q) = \int p_0 \log(p_0/q) d\mu = \int \tilde{p}_0 \log(\tilde{p}_0/\tilde{q}) dQ_0 = \mathrm{KL}_{Q_0}(\tilde{p}_0, \tilde{q}).$$

By Lusin's theorem we have a compact $E \subset \Omega$ with $Q_0(E^c) < \epsilon'$ such that $\tilde{p}_0$ is continuous on E. This E can be chosen such that for every $\epsilon > 0$, there exists a partition, $\Omega = \cup_i A_i$ with all $A_i \in \mathcal{A}^{(k)}$ for some k, such that the diameter of each $A_i \cap E$ is less than ϵ. We define

$$\delta_E(\epsilon) = \sup_{x,y \in E: |x-y|<\epsilon} |\tilde{p}_0(x) - \tilde{p}_0(y)| \quad \text{and} \quad d_i = \max\left(\sup_{A_i \cap E} \tilde{p}_0(x) + \delta_E(\epsilon), \epsilon' \right)$$

and let $D_\epsilon(\tilde{p}_0)$ be the collection of step functions $g(x) = \sum_i g_i \mathbf{1}_{A_i}(x)$ with $d_i \le g_i < d_i + \delta_E(\epsilon)$. For every $g \in D_\epsilon(\tilde{p}_0)$, let $\tilde{g}$ be the normalized version of g, that is $\tilde{g} = g / \int g dQ_0$. Then

$$\int_E (g - \tilde{p}_0) dQ_0 - \int_{E^c} |g - \tilde{p}_0| dQ_0 \le \int (g - \tilde{p}_0) dQ_0 \le \int_E (g - \tilde{p}_0) dQ_0 + \int_{E^c} |g - \tilde{p}_0| dQ_0,$$

and so

$$\delta_E(\epsilon) - (2M + \epsilon')\epsilon' \le \int (g - \tilde{p}_0) dQ_0 \le 3\delta_E(\epsilon) + (2M + \epsilon')\epsilon'.$$

Thus for any fixed ϵ, when ϵ' is small enough, we have $\int (g - \tilde{p}_0) dQ_0 \ge 0$, and thus,

$$\log\left(\int g dQ_0 \right) = \log\left(1 + \int (g - \tilde{p}_0) dQ_0 \right) \le 3\delta_E(\epsilon) + (2M + \epsilon')\epsilon'.$$

Now,

$$\begin{aligned} 0 \le \mathrm{KL}_{Q_0}(\tilde{p}_0, \tilde{g}) &= \int \tilde{p}_0 \log(\tilde{p}_0/\tilde{g}) dQ_0 \\ &= \int_E \tilde{p}_0 \log(\tilde{p}_0/g) dQ_0 + \int_{E^c} \tilde{p}_0 \log(\tilde{p}_0/g) dQ_0 + \log\left(\int g dQ_0 \right) \\ &\le M \log(M/\epsilon')\epsilon' + 3\delta_E(\epsilon) + (2M + \epsilon')\epsilon'. \end{aligned}$$

By first choosing ϵ' and then ϵ small enough, we can make $\mathrm{KL}_{Q_0}(\tilde{p}_0, \tilde{g})$ arbitrarily small. So p_0 lies in the K-L support of π. Therefore, by Schwartz's theorem, we have posterior consistency at p_0 under the weak topology. □

Proof of Lemma 2. If $n(A) = 0$ then by definition $\xi_A(i, \boldsymbol{\phi}) = 1$. If A has no children, then also by definition $\xi_A(i, \boldsymbol{\phi}) = q_0(\boldsymbol{x}|A)$. If $n(A) = 1$, then $\xi_A(i, \boldsymbol{\phi})$ becomes the conditional prior predictive density on A valued at x, which is just $q_0(x|A)$ since the Markov-APT conditional A is still an Markov-tree and by Theorem 3 its predictive density is Q_0. Finally we consider the case when A has children and $n(A) \geq 2$. For $A \in \mathcal{A}^{(\infty)} \backslash \Omega$,

$$
\begin{aligned}
\xi_A(i, \boldsymbol{\phi}) &= \int q(\boldsymbol{x}|A)\pi(dq \,|\, \boldsymbol{\phi}, C(A_p) = i) \\
&= \int q(\boldsymbol{x}|A)\pi(dq \,|\, \boldsymbol{\phi}, C(A) = i')\gamma_{i,i'}(A) \\
&= \sum_{i'=1}^{I} \int q(\boldsymbol{x} \,|\, A)\pi(dq \,|\, \boldsymbol{\phi}, C(A) = i')\gamma_{i,i'}(A) \\
&= \sum_{i'=1}^{I} \gamma_{i,i'}(A) \int \theta(A)^{n(A_l)}(1 - \theta(A))^{n(A_r)}\pi(\theta(A) \,|\, C(A) = i') \times \\
&\qquad\qquad \int q(\boldsymbol{x} \,|\, A_l)q(\boldsymbol{x} \,|\, A_r)\pi(dq \,|\, \boldsymbol{\phi}, C(A) = i') \\
&= \sum_{i'=1}^{I} \gamma_{i,i'}(A) M_A^{i'}(\boldsymbol{\theta}_0)\xi_{A_l}(i', \boldsymbol{\phi})\xi_{A_r}(i', \boldsymbol{\phi}).
\end{aligned}
$$

For $A = \Omega$, the derivation follows similarly with $\gamma_{i,i'}(A)$ replaced by $\gamma_{i'}(A)$. □

Proof of Theorem 7. This theorem follows by two applications of Bayes rule. For $A \in$

$\mathcal{A}^{(\infty)} \backslash \Omega$, the posterior transition probability is

$$\begin{aligned}
\tilde{\gamma}_{i,i'}(A) &= P(C(A) = i' \,|\, C(A_p) = i, \boldsymbol{x}(A)) \\
&= \int q(\boldsymbol{x} \,|\, A) \gamma_{i,i'}(A) \pi(dq \,|\, C(A) = i') \Big/ \int q(\boldsymbol{x} \,|\, A) \pi(dq \,|\, C(A_p) = i) \\
&= \begin{cases} \gamma_{i,i'}(A) M_A^{i'}(\boldsymbol{\theta}_0) \xi_{A_l}(i', \boldsymbol{\phi}) \xi_{A_r}(i', \boldsymbol{\phi}) / \xi_A(i, \boldsymbol{\phi}) & \text{if } A \text{ has children} \\ \gamma_{i,i'}(A) M_A^{i'}(\boldsymbol{\theta}_0) / \xi_A(i, \boldsymbol{\phi}) & \text{otherwise.} \end{cases}
\end{aligned}$$

Therefore the state transition probability matrix is

$$\tilde{\boldsymbol{\gamma}}(A) = \boldsymbol{D}'(A)^{-1} \boldsymbol{\gamma}(A) \boldsymbol{D}''(A).$$

For $A = \Omega$, the proof for the initial state probability vector is similar. The expression follows because the overall marginal likelihood is $\xi_\Omega(1, \boldsymbol{\phi})$. □

S2. Prior specification of the DPM of normals

We use the `DPpackage` function `DPdensity` to carry out density estimation using the Dirichlet process mixture (DPM) of normals. The model formulation and the hyperparameter values

are as follows, which follows an example in the user manual for `DPpackage`

$$
\begin{aligned}
x_i \,|\, \mu_i, \Sigma_i &\sim \mathrm{N}(\mu_i, \Sigma_i) \quad \text{for } i = 1, 2, \ldots, n \\
(\mu_i, \Sigma_i) \,|\, H &\sim H \\
H \,|\, \alpha, H_0 &\sim \mathrm{DP}(\alpha H_0) \\
H_0 &= \mathrm{N}(\mu | m_1, \Sigma/k_0) \times \mathrm{IW}(\Sigma | \nu_1, \psi_1) \\
\alpha \,|\, a_0, b_0 &\sim \mathrm{Gamma}(a_0, b_0) \\
m_1 \,|\, m_2, s_2 &\sim \mathrm{N}(m_2, s_2) \\
k_0 \,|\, \tau_1, \tau_2 &\sim \mathrm{Gamma}(\tau_1/2, \tau_2/2) \\
\psi_1 \,|\, \nu_2, \psi_2 &\sim \mathrm{IW}(\nu_2, \psi_2)
\end{aligned}
$$

where $a_0 = 2$, $b_0 = 1$, $m_2 = 0$, $s_2 = 10^5$, $\psi_2 = \mathrm{diagonal}(0.5, 1)$, $\nu_1 = 4$, $\nu_2 = 4$, $\tau_1 = 1$, and $\tau_2 = 100$. We draw 5,000 posterior samples using 1,000 burn-in iterations and a 10-iteration thinning window.